arXiv:1606.05277v1 [stat.ME] 16 Jun 2016

Stochastic blockmodels for exchangeable collections of networks

Perla E. Reyes and Abel Rodríguez *
Department of Applied Mathematics and Statistitics
University of California, Santa Cruz

Abstract

We construct a novel class of stochastic blockmodels using Bayesian nonparametric mixtures. These model allows us to jointly estimate the structure of multiple networks and explicitly compare the community structures underlying them, while allowing us to capture realistic properties of the underlying networks. Inference is carried out using MCMC algorithms that incorporates sequentially allocated split-merge steps to improve mixing. The models are illustrated using a simulation study and a variety of real-life examples.

1 Introduction

Network data consists of measurements associated with the interactions among a set of nodes (which we call actors), and is often visualized in the form of a (weighted) graph. Network data has become quite ubiquitous in fields as diverse as sociology, bioinformatics, finance and physics. In fact, it is often the case that multiple relationships are observed among a given set of actors, in which case it is of interest to model them jointly. Indeed, since the actors are the same, we might expect similarities

*Perla E. Reyes is Postdoctoral Associate, Department of Applied Mathematics and Statistics, University of California, Mailstop SOE2, Santa Cruz, CA 95064, perla@soe.ucsc.edu. Abel Rodriguez is Associate Professor, Department of Applied Mathematics and Statistics, University of California, Mailstop SOE2, Santa Cruz, CA 95064, abel@soe.ucsc.edu. This material is based upon work supported by DARPA Award No. N66001-10-1-4016. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of DARPA.

across the networks, which means that performing independent analyses on each network is potentially wasteful. On the other hand, since different networks might reflect slightly different modes of interaction, just collapsing all observations into a single "concensus" network, or assuming that all networks arise from the same underlying stochastic process, might yield misleading results.

Popular statistical approaches for modeling networks include the class of exponentially weighted random graphs of Frank & Strauss (1986), the class of p_1 models introduced in Holland & Leinhardt (1975), and the latent social space models of Hoff et al. (2002) and Handcock et al. (2007). This paper is concerned with building hierarchical Bayesian models for an exchangeable collection of networks, with a particular emphasis on procedures that allow for network comparison. The methods we discuss build on the popular literature on stochastic blockmodels (White et al., 1976; Wang & Wong, 1987). Stochastic blockmodels extend the notion of model-based clustering to network data. More specifically, stochastic blockmodels aim at discovering an "optimal partition" of the actors in the network into homogenous groups (the factions or communities), which are made of actors that are (approximately) structurally equivalent (Lorrain & White, 1971; Wasserman & Faust, 1994). These communities represent "social positions" or "social roles" ; members of the same faction are "substitutable", in the sense that they are subject to similar opportunities and constraints. Hence, stochastic blockmodels are appealing because of their interpretability: factions are meaningful social constructs that are often driven by unobservable (or unobserved) variables. Of particular interest to us are the subclass of infinite-dimensional blockmodels introduced by Kemp et al. (2006) and Xu et al. (2006), which treat the number of communities as an unknown parameter and place a probability distribution over all possible partitions of the set of actors.

As a motivating example, consider the study reported by Krackhardt (1987), who collected cognitive social structure data from 21 management personnel in a high-tech, machine manufacturing firm to asses the effects of a recent management intervention program. Each person indicated not only who he or she believes is their friend, but also his or her perception of others friendships; the result is a set of twenty one networks, each one of size 21×21. In this example it is natural to ask how similar the perception of the network is among the subjects involved; the answer to this question provides important insights about the social structures within network (see Krackhardt, 1987 for further discussion). In addition, once we have identified groups of individuals with similar perceptions, it would be natural to try to aggregate information within each of these groups to improve estimation of the underlying parameters determining the network structure. Another motivating example comes from the study on the interactions among 14 employees in a Western Electric plant reported by Roethlisberger & Dickson (1939). The researchers recorded six types of interactions among the employees: friendship,

participation in horseplay, helping others with work, antagonistic behavior, arguments about open windows and number of times workers traded job assignments. Again, it would be natural to compare the social structure associated with the different relationships; however, this task is particularly challenging because of the different data types involved in the analysis (which includes directed and undirected networks, as well as binary and count data).

Recently, the literature on blockmodels, and more generally, community identification algorithms, has been very active. Girvan & Newman (2002), Newman (2004), Clauset (2005) and Mishra et al. (2008) are some recent examples within the physics and machine learning literature. However, most of these approaches are algorithmic and it is unclear how uncertainty estimates can be obtained. On the statistics side, two recent publications are particularly noteworthy. On one hand, Airoldi et al. (2008) develop mixed-membership blockmodels, where subjects can be members of more than one community simultaneously. On the other, Bickel & Chen (2009) show that likelihood-based community detection algorithms are asymptotically consistent, while algorithms based on certain modularity scores (such as the one described in Girvan & Newman, 2002) are not. In any case, all the approaches discussed above deal with single network problems rather than with the analysis of multiple networks, which is the focus of this manuscript.

Indeed, although network comparison and aggregation are issues that seem to appear often in practice, we are not aware of work on model-based statistical methods to address them. Contingency tables have been used for modeling multi-relational data at least since Galaskiewicz & Marsden (1978). However, this type of "macroanalysis" focuses solely on the relations and ignores individual actors. Fienberg et al. (1985) extend the class of p_1 models to multi-relational data. However, their approach relies on a partition of the actors into subgroups that is generated from external (extra-relational) information, which is assumed to be common to all relationships. Moreover, these approaches are restricted binary networks and cannot be applied to problems with mixed data types.

The dearth of formal statistical procedures for the comparison of network structures might be explained in part by the lack of standard asymptotic results. To ascertain the similarity among networks, practitioners usually contrast summary statistics (such as indegree, outdegree, or betweenness distributions) across networks using simple tests. However, this approach is often inappropriate, specially when networks correspond to different data types and there is no obvious way to construct a common summary that applies across all networks (as is the case, for example, of the Western Electric example mentioned before). In the case of aggregation, commonly employed methods assign a "consensus" value to the interaction between two subjects across several networks (for an example, see Krackhardt, 1987). A typical aggregate would be a weighted average of all observed networks. An obvious

drawback of this approach is that choosing the weights is not straightforward, and assigning the same weight to all networks will often be inappropriate. Furthermore, this type of approach cannot be applied if networks correspond to different data types.

One important feature of the models developed in this paper is their ability to account for complex features of the networks such as assortative / disassortative mixing, and to incorporate prior information about the degree distribution. Indeed, existing approaches to Bayesian inference in stochastic blockmodels pay little attention to prior elicitation and the effect of the prior on inference. In addition to discussing models for multiple networks, this paper explores the *a priori* properties of networks generated by stochastic blockmodels and discusses hierarchical specifications that can be used to include prior information about the topology of the network.

The remaining of the paper is organized as follows: Section 2 describes a general framework for modeling a single network and reviews the use of nonparametric mixture priors in the context of network models. Section 3 presents generalization of this class of models that allow for more flexible prior specification. Section 4 extends the single-network model to a multiple networks model that simultaneously identifies community structure per network and establishes similarities across networks. This section also discusses the MCMC algorithm that we use for posterior inference. Key features of the multiple networks model are showcase in Section 5 using first a simulation study and then datasets of employee relationships. Finally, Section 6 presents our conclusions.

2 Stochastic blockmodels

For the purpose of this paper, a network is a $I \times I$ matrix $\mathbf{Y} = [y_{i,j}]$, where $y_{i,j}$ measures the strength of the relationship between actor i and actor j. The network is called undirected if $\mathbf{Y}$ is symmetric (e.g., if it is irrelevant who initiates the interaction), and directed otherwise. Further, the network is called acyclic if subjects do not interact with themselves, in which case the diagonal elements of $\mathbf{Y}$ are treated as structural zeros[1].

Stochastic blockmodels (White et al., 1976; Wang & Wong, 1987) extend the notion of model-based clustering to networks. In the case of acyclic, directed networks, Bayesian stochastic blockmodels are hierarchical models that for $i, j = 1, \ldots, I$ and $j \neq i$.

$$y_{i,j}|\xi_i, \xi_j, \boldsymbol{\Theta} \sim \psi(y_{ij}|\theta_{\xi_i,\xi_j}, \boldsymbol{\nu}), \quad \xi_i|\mathbf{w} \sim \sum_{k=1}^{N} w_k \delta_k, \quad (\mathbf{w}, \boldsymbol{\Theta}, \boldsymbol{\nu}|\boldsymbol{\varsigma}, \boldsymbol{\lambda}) \sim p(\mathbf{w}|\boldsymbol{\varsigma}) p(\boldsymbol{\Theta}|\boldsymbol{\lambda}) p(\boldsymbol{\nu}), \quad (1)$$

[1] We deviate slightly from the standard definition of acyclic networks, which typically applies only to directed networks and precludes any sort closed loops.

where N is the maximum number of potential factions, $\boldsymbol{\varsigma}$ and $\boldsymbol{\lambda}$ are vectors of hyperparameters (which will be assigned hyperpriors $p(\boldsymbol{\varsigma})$ and $p(\boldsymbol{\lambda})$), $\boldsymbol{\Theta} = [\theta_{k,l}]$ is a $N \times N$ matrix, $\mathbf{w} = (w_1, \ldots, w_N)'$ is such that $\sum_{k=1}^{N} w_k = 1$, δ_k denotes the degenerate probability distribution placing probability 1 on k, and $\psi(y|\theta, \boldsymbol{\nu})$ is parametric kernel indexed by the parameters θ and $\boldsymbol{\nu}$, where θ is a random effect, and $\boldsymbol{\nu}$ is a vector of fixed effects. In the case of undirected networks, a similar definition applies with the added constrains $y_{i,j} = y_{j,i}$ and $\theta_{k,l} = \theta_{l,k}$.

The formulation in (1) is extremely flexible and easily interpretable. The latent variables $\boldsymbol{\xi} = (\xi_1, \ldots, \xi_I)'$ act as (unobserved) faction indicators; the prior probability that any two subjects are assigned to the same cluster (i.e., the share the same social role) is given by $\mathsf{E}\left\{\sum_{k=1}^{N} w_k^2\right\}$. Binary networks (i.e., those where $y_{i,j} \in \{0, 1\}$, so that $y_{i,j} = 1$ if actor i interacts with actor j, and $y_{i,j} = 0$ otherwise) can be accommodated by taking $y_{i,j}|\xi_i, \xi_j, \boldsymbol{\Theta} \sim \mathsf{Ber}(\theta_{\xi_i,\xi_j})$ and selecting (for computational convenience) $\theta_{k,l} \sim \mathsf{Beta}(a, b)$. In this case, the entries $\theta_{k,l}$ give the probability that an interaction occurs between actor from factions k and l. On the other hand, count data could be incorporated by taking $y_{i,j}|\xi_i, \xi_j, \boldsymbol{\Theta} \sim \mathsf{Poi}(\theta_{\xi_i,\xi_j})$ and $\theta_{k,l} \sim \mathsf{Gam}(a, b)$. In this case, $\theta_{k,l}$ is the intensity of the interaction between factions k and l.

One example of a stochastic blockmodel is the infinite relational blockmodel (IRM) (Kemp et al., 2006; Xu et al., 2006). In the IRM, $N = \infty$ and

$$\theta_{k,l} \sim_{iid} H^{\boldsymbol{\lambda}}, \qquad w_k = v_k \prod_{s<k}(1 - v_s), \qquad v_k \sim_{iid} \mathsf{Beta}(1, \beta), \qquad k, l = 1, 2, \ldots, \tag{2}$$

where $H^{\boldsymbol{\lambda}}$ is a parametric distribution indexed by the hyperparmeter $\boldsymbol{\lambda}$. This structure for the weights, which is strongly connected to the stick-breaking construction of the Dirichlet process (Sethuraman, 1994), implies that the joint distribution of $\boldsymbol{\xi} = (\xi_1, \ldots, \xi_I)$ obtained *after* integrating out $\mathbf{w}$ can be described by a sequence of predictive distributions with $\xi_1 = 1$ and

$$\xi_i|\xi_{i-1}, \ldots, \xi_1 \sim \sum_{k=1}^{K^{i-1}} \frac{m_k^{i-1}}{\beta + i - 1}\delta_k + \frac{\beta}{\beta + i - 1}\delta_{K^{i-1}+1}, \qquad 2 \leq i \leq I, \tag{3}$$

where δ_a denotes the degenerate probability distribution on a, $K^{i-1} = \max_{j<i}\{\xi_j\}$ is the number of unique values among $\xi_1, \ldots, \xi_{i-1}$, $m_k^{i-1} = \sum_{j=1}^{i-1} \mathbf{1}_{(\xi_j = k)}$ is the number of indicators equal to k among $\xi_1, \ldots, \xi_{i-1}$, and $\beta > 0$ is a constant. This sequence of predictive distributions is sometimes called the Chinese restaurant process (CRP) (see Figure 1), and implies that $\mathsf{Pr}(\xi_i = \xi_j) = \sum_{k=1}^{\infty} \mathsf{E}\{w_k^2\} = 1/(1 + \beta)$ for all i and j.

The CRP places a probability distribution on all possible partitions of I actors, whose shape is controlled by the parameter β. For example, the probability that all actors are assigned to a single

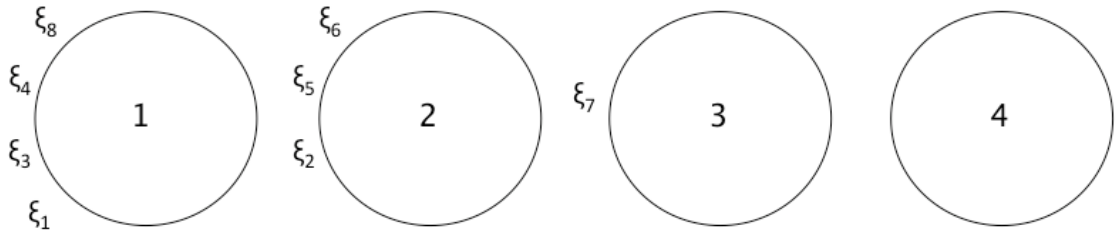

Figure 1: Schematic representation for the Chinesse restaurant process. In this analogy, observations correspond to customers who plan to dine in a restaurant with an infinite number of tables. The variable ξ_i indicates which table is occupied by customer i. Customer 1 sits at table 1; customer i sits at any of the occupied tables with probability proportional to the number of customers already sitting in it, while she sits on the next empty table with probability proportional to β. In this particular example, the next customer (number 9) would sit on table 1 with probability $4/(\beta+8)$, on table 2 with probability $3/(\beta+8)$, on table 3 with probability $1/(\beta+8)$, or on table 4 with probability $\beta/(\beta+8)$.

faction is $\prod_{i=2}^{I} \frac{i-1}{\beta+i-1}$, in which case the model reduces to an Erdös-Rényi random graph (Erdös & Rényi, 1959) with an unknown link frequency. On the other hand, the probability that each actor is assigned to a different faction is given by $\prod_{i=2}^{I} \frac{\beta}{\beta+i-1}$. Hence, although the number of potential factions N is infinite, the effective number of factions K actually occupied by actors is treated as a random variable taking values with support on the set $\{1, \ldots, I\}$, and is automatically estimated from the data. This prior distribution for K implied by the CRP is such that $\mathsf{E}(K) \sim \beta \log\{\frac{\beta+I}{\beta}\}$ for large I.

3 Marginal model properties and prior elicitation

The IRM described in the previous Section has two unappealing properties. Firstly, note that for a binary network under (1) and (2), $\mathsf{Pr}(y_{i,j} = 1|\xi_i = \xi_j) = \mathsf{Pr}(y_{i,j} = 1|\xi_i \neq \xi_j)$ a priori, i.e., membership in the same faction does not provide any information about how likely a link is. This flies in the face of well known empirical facts (for example, most social networks present assortative mixing, where members of the same faction have a higher probability of interacting). Secondly, note that $\mathsf{Pr}(y_{i,j}|\boldsymbol{\lambda}) = \mathsf{E}_{H^{\boldsymbol{\lambda}}}(\theta_{k,l}) = \bar{\theta}$ for every β. Moreover, the expected number of links per actor is simply given by $(I-1)\mathsf{E}_{H^{\boldsymbol{\lambda}}}(\theta_{k,l})$. Hence, prior information about the community structure provides

no information about the expected degree in the network.

To alleviate these issues we consider a more general prior specification such that

$$p(\boldsymbol{\Theta}|\boldsymbol{\lambda}=(\boldsymbol{\lambda}_O,\boldsymbol{\lambda}_D)) = \begin{cases} \left\{\prod_{k=1}^{\infty} H^{\boldsymbol{\lambda}_D}(\theta_{k,k})\right\}\left\{\prod_{k=1}^{\infty}\prod_{l=k+1}^{\infty} H^{\boldsymbol{\lambda}_O}(\theta_{k,l})\delta_{\theta_{k,l}}(\theta_{l,k})\right\} & \mathbf{Y} \text{ undirected} \\ \left\{\prod_{k=1}^{\infty} H^{\boldsymbol{\lambda}_D}(\theta_{k,k})\right\}\left\{\prod_{k=1}^{\infty}\prod_{l=k+1}^{\infty} H^{\boldsymbol{\lambda}_O}(\theta_{k,l})H^{\boldsymbol{\lambda}_O}(\theta_{l,k})\right\} & \mathbf{Y} \text{ directed} \end{cases}, \tag{4}$$

where $H^{\boldsymbol{\lambda}}$ is a distribution function indexed by a parameter $\boldsymbol{\lambda}$, and $p(\mathbf{w})$ is defined by

$$w_k = u_k \prod_{s<k}(1-u_s), \qquad u_k \sim_{iid} \mathsf{Beta}(1-\alpha, \beta+\alpha k), \qquad k = 1, 2, \ldots. \tag{5}$$

This Section explores the properties of this new specification, with a particular emphasis on 1) studying the *a-priori* properties of the network, 2) determining procedures for prior specification, and 3) generating model-based alternatives to traditional summaries of network topology. Most of our discussion is focused on binary networks because they are arguably the most common type of networks found in practice. However, most of the comments can be easily extended to more general classes of networks.

3.1 Exchangeability of actors and communities

An infinite two-dimensional array $\mathbf{Y}$ is said to be jointly exchangeable if, for any finite submatrix, the distribution is unchanged when the same permutation is applied to both the rows and the columns of the submatrix. In the absence of node-specific covariate information, this is often a natural assumption for networks that simply implies that the order in which the actors are observed should have to influence in our probability model.

The blockmodel defined in Equation (1) defines a jointly infinitely exchangeable prior on inifnite arrays; this can be most easily seen by exploiting the representation theorem from Aldous (1981): an array is jointly exchangeable if we can write $y_{i,j} = g(\boldsymbol{\phi}, v_i, v_j, u_{i,j})$ where $\boldsymbol{\phi}$, $\{v_i\}$ and $u_{i,j}$ are independent random variables such that $v_i \sim \mathsf{Uni}[0,1]$ and $u_{i,j} \sim \mathsf{Uni}[0,1]$. In the case of blockmodels, we can take $\boldsymbol{\phi} = (\mathbf{w}, \boldsymbol{\Theta})$ and write

$$y_{i,j} = g(\boldsymbol{\phi}, v_i, v_j, u_{i,j}) = \begin{cases} 1 & u_{i,j} \leq \theta_{G^{-1}(v_i),G^{-1}(v_j)} \\ 0 & u_{i,j} > \theta_{G^{-1}(v_i),G^{-1}(v_j)} \end{cases},$$

where $G(\cdot) = \sum_k w_k \delta_k$ and $G^{-1}(\cdot)$ is the generalized inverse of G. Hence, the representation theorem ensures that the actors are exchangeable, no matter what the priors on $\mathbf{w}$ and $\boldsymbol{\Theta}$ are.

Although ensuring exchangeability in the matrix $\mathbf{Y}$ does not impose constrains on our choice of priors for $\boldsymbol{\Theta}$ and $\mathbf{w}$, in practice we are often interested in models where not only the individuals but also the unknown communities are exchangeable. In order to ensure that a blockmodel induces exchangeability among factions, we need the prior on $\boldsymbol{\Theta}$ to also be jointly exchangeable. The following lemma ensures that our generalized model satisfies this constraint.

Lemma 1 *Let $p(\boldsymbol{\Theta}|\boldsymbol{\lambda})$ be defined as in* (4)*, and let $\boldsymbol{\lambda} \sim p(\boldsymbol{\lambda})$. Then, $p(\boldsymbol{\Theta}) = \int p(\boldsymbol{\Theta}|\boldsymbol{\lambda})p(\boldsymbol{\lambda})d\boldsymbol{\lambda}$ defines a jointly exchangeable prior on $\boldsymbol{\Theta}$.* ◇

The proof can be found in Appendix A. One implication of Equation (11) is that the a priori marginal probability of a link is given by a weighted average of the probability of a link under $H^{\boldsymbol{\lambda}_D}$ and $H^{\boldsymbol{\lambda}_O}$.

Lemma 2 *Let $\mathbf{Y}$ be an undirected network. For the blockmodel in* (1) *and the priors in* (4) *and* (5)*, the a priori marginal probability of a link between any two nodes i and j is given by*

$$\bar{\theta} = \mathsf{Pr}(y_{ij} = 1|\boldsymbol{\lambda}, \boldsymbol{\varsigma}) = \left(\frac{1-\alpha}{\beta+1}\right) \mathsf{E}_{H^{\boldsymbol{\lambda}_D}} \{\theta_{k,k}\} + \left(\frac{\beta+\alpha}{\beta+1}\right) \mathsf{E}_{H^{\boldsymbol{\lambda}_O}} \{\theta_{k,l}\}.$$

Under the same circumstances, but for a directed network $\mathbf{Y}$, we have $\bar{\theta}^{in} = \bar{\theta}^{out} = \bar{\theta}$, where $\bar{\theta}^{in} = \mathsf{Pr}(y_{ji} = 1|\boldsymbol{\lambda}, \boldsymbol{\varsigma})$ and $\bar{\theta}^{out} = \mathsf{Pr}(y_{ij} = 1|\boldsymbol{\lambda}, \boldsymbol{\varsigma})$ represent the prior probabilities of a link from j to i and a link from i to j respectively. ◇

For a proof, see Appendix B. Note that $(1-\alpha)/(\beta+1)$ represents the percentage of links of node i with members of its same community, and

$$\min\{\mathsf{E}_{H^{\boldsymbol{\lambda}_D}}\{\theta_{k,k}\}, \mathsf{E}_{H^{\boldsymbol{\lambda}_O}}\{\theta_{k,l}\}\} \leq \bar{\theta} \leq \max\{\mathsf{E}_{H^{\boldsymbol{\lambda}_D}}\{\theta_{k,k}\}, \mathsf{E}_{H^{\boldsymbol{\lambda}_O}}\{\theta_{k,l}\}\}$$

for any choice of prior on $\mathbf{w}$. On the other hand, if $\boldsymbol{\lambda}_D = \boldsymbol{\lambda}_O$ we recover $\mathsf{Pr}(y_{ij} = 1|\boldsymbol{\lambda}, \boldsymbol{\varsigma}) = \mathsf{E}_{H^{\boldsymbol{\lambda}}}(\theta_{k,l})$ as in traditional IRMs.

3.2 Capturing assortative / disassortative mixing

A network is said to be assortative (or have assortative mixing) if interactions among actors in the same faction tend to be more common than interactions among actors in different factions. Similarly, a network is said to be disassortative if the opposite is true. It is often the case that we have a-priori knowledge about the mixing pattern associated with a specific network; for example, social networks are often assortative, while ecological prey-predator networks are often disassortative.

One advantage of (4) is that, by allowing different hyperparameters for the diagonal and off-diagonal elements of $\boldsymbol{\Theta}$, it allows us to incorporate information about the type of mixing in the network. For example, for an undirected binary network with $y_{ij}|\xi_i, \xi_j, \boldsymbol{\Theta} \sim \mathsf{Ber}(\theta_{\xi_i,\xi_j})$ which is believed to have assortative mixing, we might assign $\theta_{k,k} \sim \mathsf{Beta}(a_D, b_D)$ and $\theta_{k,l} \sim \mathsf{Beta}(a_O, b_O)$, where $\boldsymbol{\lambda}_D = (a_D, b_D)$ and $\boldsymbol{\lambda}_O = (a_O, b_O)$ are given a joint prior $p(a_D, b_D, a_O, b_O)$ such that $\{a_D/(a_D+b_D) > a_O/(a_O + b_O)\}$ has high probability.

This discussion suggests that a simple way to summarize the effect of our prior choice on the topology of the network is through the assortativity index

$$\Upsilon = \log\left(\mathsf{E}_{H^{\boldsymbol{\lambda}_D}}\{\theta_{k,k}|\boldsymbol{\lambda}_D\}\right) - \log\left(\mathsf{E}_{H^{\boldsymbol{\lambda}_O}}\{\theta_{k,l}|\boldsymbol{\lambda}_O\}\right).$$

Note that $p(\boldsymbol{\lambda}_D) = p(\boldsymbol{\lambda}_O)$ implies $\mathsf{E}\{\Upsilon\} = 0$ a priori, which means that a priori we have no information about whether the network is assortative or disassortative; $\mathsf{E}\{\Upsilon\} > 0$ is associated with assortative networks, while $\mathsf{E}\{\Upsilon\} < 0$ is associated with disassortative networks. A posteriori, the distribution of Υ provides a simple summary of the type of mixing in the network and a model-based alternative to the assortativity coefficients discussed in Newman (2003a).

3.3 Degree distribution

The degree of a node refers to the number of links associated with it. Hence, for undirected binary networks, the degree of actor i is simply $D_i = \sum_{j \neq i} y_{ij}$. For directed networks, we can analogously define the in-degree D_i^I and the out-degree D_i^O of actor i as the number of links that start or end at actor i, $D_i^I = \sum_{j \neq i} y_{ji}$ and $D_i^O = \sum_{j \neq i} y_{ij}$. The distribution of D_i is often used to compare how well analytical models fit the observed data.

Lemma 3 *Consider an undirected network* $\mathbf{Y}$. *For the blockmodel in* (1) *and the priors in* (4) *and* (5)*, the degree distribution satisfies*

$$\begin{aligned}
\mathsf{E}\{D_i|\boldsymbol{\lambda}, \boldsymbol{\varsigma}\} = \bar{\rho} &= [I-1]\left[\left(\frac{1-\alpha}{\beta+1}\right)\mathsf{E}_{H^{\boldsymbol{\lambda}_D}}\{\theta_{k,k}\} + \left(\frac{\beta+\alpha}{\beta+1}\right)\mathsf{E}_{H^{\boldsymbol{\lambda}_O}}\{\theta_{k,l}\}\right] \\
\mathsf{Var}(D_i|\boldsymbol{\lambda}, \boldsymbol{\varsigma}) = \bar{\kappa} &= \bar{\rho}(1-\bar{\rho}) + \frac{[I-1][I-2]}{2(\beta+1)(\beta+2)}\Big[(1-\alpha)(2-\alpha)\mathsf{E}_{H^{\boldsymbol{\lambda}_D}}\{\theta_{k,k}^2\} \\
&\quad + (\beta+\alpha)(\beta+2\alpha)\left(\mathsf{E}_{H^{\boldsymbol{\lambda}_O}}\{\theta_{k,l}\}\right)^2 \\
&\quad + (1-\alpha)(\beta+\alpha)\left[2\mathsf{E}_{H^{\boldsymbol{\lambda}_D}}\{\theta_{k,k}\}\mathsf{E}_{H^{\boldsymbol{\lambda}_O}}\{\theta_{k,l}\} + \mathsf{E}_{H^{\boldsymbol{\lambda}_O}}\{\theta_{k,l}^2\}\right]\Big],
\end{aligned}$$

where I is the number of actors in the network. If $\mathbf{Y}$ is a directed network, a similar result holds separately for the in-degree and the out-degree, and we also have

$$\mathsf{Cor}(D_i^{in}, D_i^{out}|\boldsymbol{\lambda}, \boldsymbol{\varsigma}) = \mathsf{Cor}(y_{ij}, y_{ji}|\boldsymbol{\lambda}, \boldsymbol{\varsigma}) = \bar{\Delta} = \frac{\bar{\Delta}^N}{\bar{\Delta}^D},$$

where

$$\bar{\Delta}^N = \left(\frac{1-\alpha}{\beta+1}\right)\mathsf{E}_{H^{\boldsymbol{\lambda}_D}}\{\theta_{k,k}^2\} + \left(\frac{\beta+\alpha}{\beta+1}\right)\left(\mathsf{E}_{H^{\boldsymbol{\lambda}_O}}\{\theta_{k,l}\}\right)^2 - \left[\left(\frac{1-\alpha}{\beta+1}\right)\mathsf{E}_{H^{\boldsymbol{\lambda}_D}}\{\theta_{k,k}\} + \left(\frac{\beta+\alpha}{\beta+1}\right)\mathsf{E}_{H^{\boldsymbol{\lambda}_O}}\{\theta_{k,l}\}\right]^2$$

and

$$\bar{\Delta}^D = \left(\frac{1-\alpha}{\beta+1}\right)\mathsf{E}_{H^{\boldsymbol{\lambda}_D}}\{\theta_{k,k}\} + \left(\frac{\beta+\alpha}{\beta+1}\right)\mathsf{E}_{H^{\boldsymbol{\lambda}_O}}\{\theta_{k,l}\} - \left[\left(\frac{1-\alpha}{\beta+1}\right)\mathsf{E}_{H^{\boldsymbol{\lambda}_D}}\{\theta_{k,k}\} + \left(\frac{\beta+\alpha}{\beta+1}\right)\mathsf{E}_{H^{\boldsymbol{\lambda}_O}}\{\theta_{k,l}\}\right]^2 = \hat{\theta}\left\{1-\hat{\theta}\right\}.$$

$\diamond$

The proof is in Appendix C. Moreover, the moment generating function for D_i can be written as (see Appendix D)

$$\mathsf{E}\left\{\left[1+\sum_{s=1}^{m_1-1}\binom{m_1-1}{s}\mathsf{E}_{H^{\boldsymbol{\lambda}_D}}(\theta_{1,1}^s)(e^t-1)^s\right]\prod_{k=2}^{K}\left[1+\sum_{s=1}^{m_k}\binom{m_k}{s}\mathsf{E}_{H^{\boldsymbol{\lambda}_O}}(\theta_{1,k}^s)(e^t-1)^s\right]\right\}, \quad (6)$$

where the last expectation is taken with respect to the joint distribution of the number of factions and faction sizes, which is given by (Pitman, 1995)

$$p(K, m_1, \ldots, m_K) = \frac{\Gamma(\beta+1)}{(\beta+K\alpha)\Gamma(\beta+I)}\prod_{k=1}^{K}(\beta+k\alpha)\frac{\Gamma(m_k-\alpha)}{\Gamma(1-\alpha)}$$

where $1 \leq K \leq I$, $1 \leq m_k \leq I$ and $\sum_{k=1}^{K} m_k = I$. Although obtaining closed-from results for the full degree distribution is challenging, our results on the mean and variance of the degree distributions, together with information about the assortative/disassortive structure in the network and/or the number of communities, can be used to elicit meaningful informative priors for network models.

To explore the effect of diverse values of α, β, and assumptions on H_D and H_O in the degree distribution of the resulting network we conducted a simulation study. For the stick-breaking parameters, we set α in $\{0, 0.2, 0.5, 0.8\}$ and β in $\{0.5, 1.5, 5.0\}$. H_D and H_O are selected such

that $(\bar{\theta}_D, \bar{\theta}_O) = (\mathsf{E}_{H^{\lambda_D}}\{\theta_{kk}\}, \mathsf{E}_{H^{\lambda_O}}\{\theta_{kl}\})$ are in $\{(0.2, 0.2), (0.2, 0.8), (0.8, 0.2), (0.8, 0.8)\}$ with 3 levels of variability: 1) point-masses at $\bar{\theta}_D$ and $\bar{\theta}_O$ (no variability), 2) beta distributions with parameters $a_D + b_D = a_O + b_O = 0.5$ (high variability), and 3) beta distributions with parameters $a_D + b_D = a_O + b_O = 5$ (low variability). For each of the different combinations of parameters 10,000 networks with $I = 100$ actors each were generated.

The resulting degree distributions are presented in Figure 2, each plot shows $\mathsf{Pr}(D_i \geq k)$ for a given $(\bar{\theta}_D, \bar{\theta}_O)$ pair and α, the first row shows results for $\alpha = 0$ which is the blocking model defined on (2), the second and third row use model (5) with $\alpha = 0.2$ (small) and $\alpha = 0.8$ (large), respectively. We observe that 1) when $\bar{\theta}_D = \bar{\theta}_O$ the specification of the cluster modeling (α and β) has no effect on the degree distribution; 2) the stick-breaking parameters have a bigger influence when assortative or disassortative behavior is present, but the specification of the prior on $\boldsymbol{\Theta}$ is still the most influential feature again.

3.4 Transitivity

Transitivity is the tendency that makes more likely the connection between two actors when both of them are related to a common third one (the friend of your friend is likely also to be your friend). Intuitively, the presence of transitivity in the network implies an increased (or decreased) number of triangles (sets of three elements all connected to each other). More formally, we can define the transitivity coefficient $\chi_{ijh} = \mathsf{Pr}(y_{ij} = 1 | y_{ih} = 1, y_{jh} = 1, \boldsymbol{\lambda}, \boldsymbol{\varsigma})$; $\chi_{ijh} > \mathsf{Pr}(y_{ij} | \boldsymbol{\lambda}, \boldsymbol{\varsigma}) = \bar{\theta}$ implies positive transitivity, $\chi_{ijh} < \bar{\theta}$ implies negative transitivity, and $\chi_{ijh} = \bar{\theta}$ implies no transitivity in the network.

Lemma 4 *For the blockmodel in* (1) *and the priors in* (4) *and* (5)*, the a priori transitivity index for a binary undirected network is given by*

$$\mathsf{Pr}(y_{ij} = 1 | y_{ih} = 1, y_{jh} = 1, \boldsymbol{\lambda}, \boldsymbol{\varsigma}) = \bar{\chi} = \bar{\chi}^N / \bar{\chi}^D,$$

where

$$\bar{\chi}^N = \left\{ \frac{(1-\alpha)(2-\alpha)}{(\beta+1)(\beta+2)} \right\} \mathsf{E}_{H^{\lambda_D}}\{\theta_{k,k}^3\} + 3 \left\{ \frac{(1-\alpha)(\beta+\alpha)}{(\beta+1)(\beta+2)} \right\} \mathsf{E}_{H^{\lambda_D}}\{\theta_{k,k}\} \mathsf{E}_{H^{\lambda_O}}\{\theta_{k,l}^2\} + \left\{ \frac{(\beta+\alpha)(\beta+2\alpha)}{(\beta+1)(\beta+2)} \right\} \left(\mathsf{E}_{H^{\lambda_O}}\{\theta_{k,l}\}\right)^3$$

and

$$\bar{\chi}^D = \left(\frac{1-\alpha}{\beta+1} \right) \mathsf{E}_{H^{\lambda_D}}\{\theta_{k,k}^2\} + \left(\frac{\beta+\alpha}{\beta+1} \right) \left(\mathsf{E}_{H^{\lambda_O}}\{\theta_{k,l}\}\right)^2$$

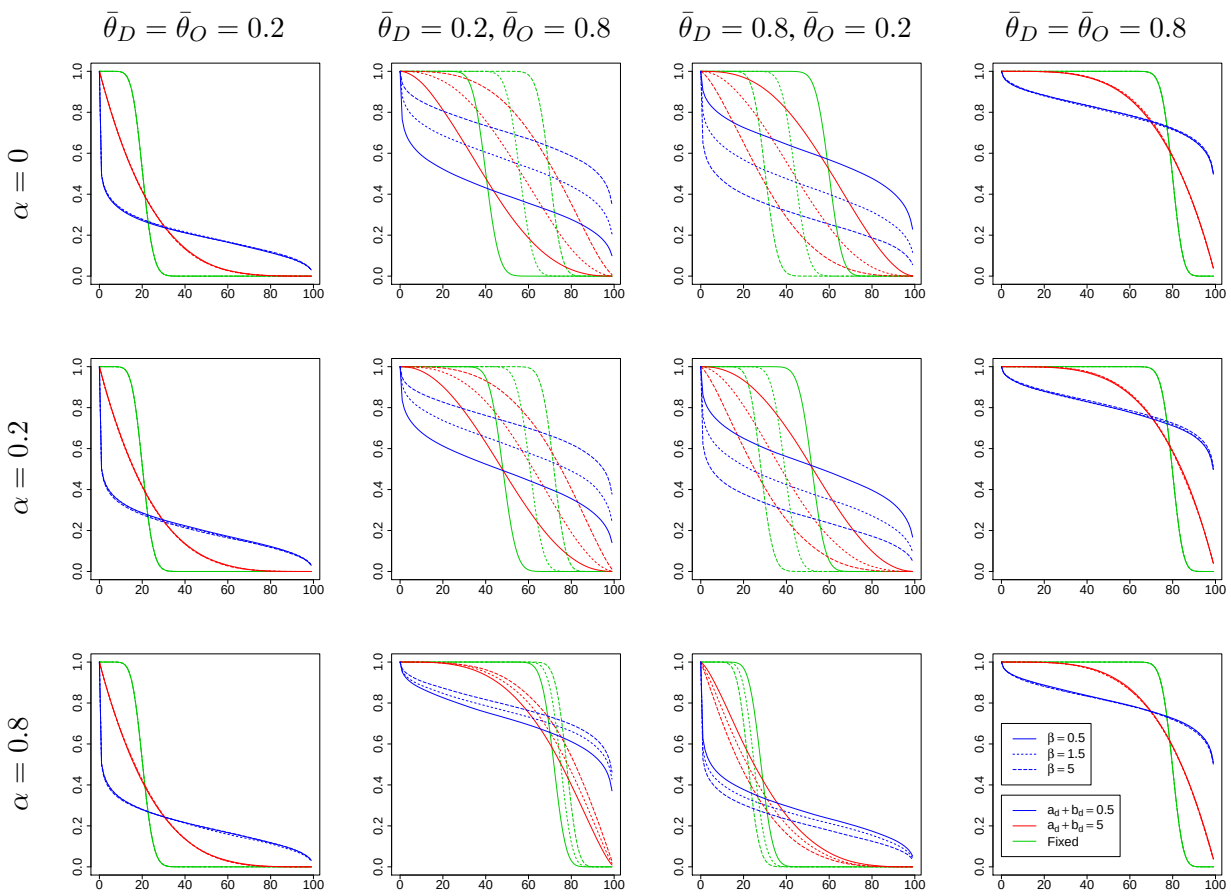

Figure 2: Degree distributions. $\Pr(D_i \geq k)$ is plotted on the vertical axis vs k on the the horizontal axes. Each line represents a combination of diverse values of the stick-breaking parameter $\beta = 0.5$ (`solid`), $\beta = 1.5$ (`dotted`) and $\beta = 5$ (`dashed`); and levels of variability, point-mass (`green`), $a_D + b_D = a_O + b_O = 5$ (`red`) and $a_D + b_D = a_O + b_O = 0.5$ (`blue`).

for any i, j and h. ◇

The proof can be seen in Appendix E. Note that $\chi_{ijk} \neq \bar{\theta}$ even under (4); hence all blockmodels assume a priori transitivity in the relationship among the subjects.

To better understand the effect of different parameters on the transitivity index we performed a second simulation study. The transitivity index is empirically approximated using the so-called clustering coefficient (Newman, 2003b),

$$C = \frac{3 \times \text{number of triangles}}{\text{number of connected triples}}$$

where a "connected triple" means a single actor that interacts with two actors. The factor of three in the numerator guaranties that $0 \leq C \leq 1$ by accounting for the fact that each triangle contributes with three triples. Using the same simulation setting than in the case of the degree distribution, 10,000 networks with $I = 100$ actors each were generated to obtain the mean (expected) clustering coefficient C. For comparison, we also found the expected value C of 10,000 networks with $I = 100$ actors each, assuming a single component model holding everything else the same, i.e. all actors belong to the same group and have probability of connection θ, where the prior $p(\theta)$ agrees with $p(\boldsymbol{\Theta}|\boldsymbol{\lambda})$. Figure 3 reports the mean clustering coefficient, on the y axis we have the (expected) clustering coefficient C, on the x axis the mean number of factions formed by the stick-breaking process (5) for the diverse values of α and β. Each row represents different levels of incertitude in the prior of $\boldsymbol{\Theta}$. In each plot, horizontal dotted lines marked the expected value of C under the single component model specification.

From the simulation results we observed that the clustering coefficient C is highly associated with $\mathsf{E}(\boldsymbol{\Theta}|\boldsymbol{\lambda})$. In fact, as with the degree distribution, $C \approx \mathsf{E}(\boldsymbol{\Theta}|\boldsymbol{\lambda})$ when $\bar{\theta}_D = \bar{\theta}_O$, regardless the specified α and β. However, C is further from $\bar{\theta}$ and always greater for small values of α, β (fewer factions) and high $\mathsf{Var}(\boldsymbol{\Theta}|\boldsymbol{\lambda})$. For the dissasortative model ($\bar{\theta}_D < \bar{\theta}_O$), C is again very close to $\mathsf{E}(\boldsymbol{\Theta}|\boldsymbol{\lambda})$, the most interesting results are obtained for the assortative model where the transitivity of the model is always larger than $\mathsf{E}(\boldsymbol{\Theta}|\boldsymbol{\lambda})$ and for small values of α, β, C could be greater than $\bar{\theta}$ by 5 - 10 percentage points. This is not the case for the single component model, whose C resulted equal to $\mathsf{E}(\theta)$ even for different levels of variability.

4 Hierarchical stochastic blockmodels for collections of networks

Consider now a situation where multiple networks are observed *for the same set of actors*, so that the data corresponds to an exchangeable collection matrices $\mathbf{Y}_1, \ldots, \mathbf{Y}_J$, where $\mathbf{Y}_j = [y_{i,i',j}]$, $i, i' = 1, \ldots, I$ and $j = 1, \ldots, J$. We are interested in jointly modeling $\mathbf{Y}_1, \ldots, \mathbf{Y}_J$ in order to improve estimation of the community structure associated with each network. Moreover, we are interested in identifying groups of networks (relationships) with similar community structure. As we discussed in the introduction, communities are meaningful constructs (e.g., in the case of social networks they can be interpreted as social roles or social positions) that are often driven by unobservable (or unobserved) variables. Hence, clusters of networks can provide important insights about the underlying processes generating the networks. Furthermore, by jointly modeling through the community structure the model can accommodate different types of networks (binary, count, directed, undirected, etc).

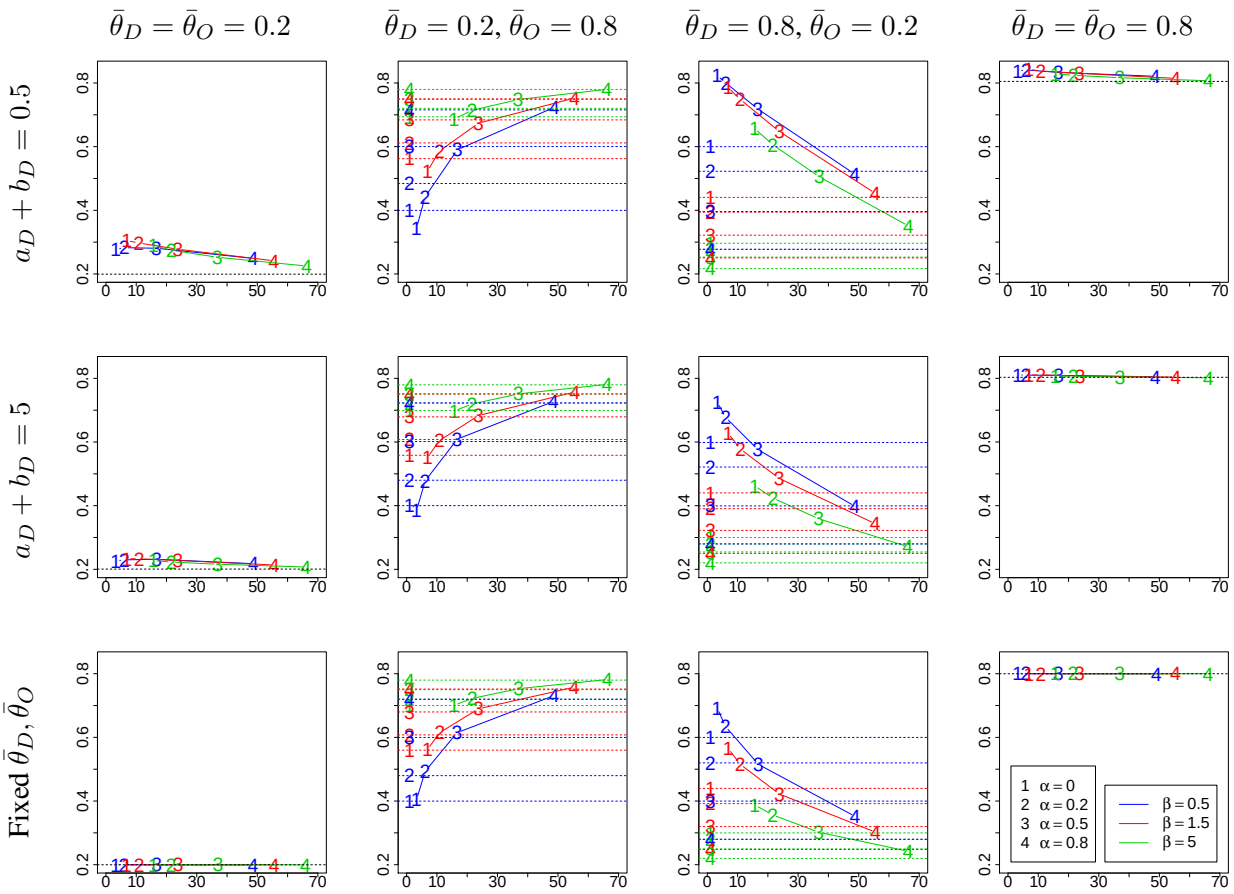
$\bar{\theta}_D = \bar{\theta}_O = 0.2$
$\bar{\theta}_D = 0.2, \bar{\theta}_O = 0.8$
$\bar{\theta}_D = 0.8, \bar{\theta}_O = 0.2$
$\bar{\theta}_D = \bar{\theta}_O = 0.8$
$a_D + b_D = 0.5$
$a_D + b_D = 5$
Fixed $\bar{\theta}_D, \bar{\theta}_O$
1 α = 0
2 α = 0.2
3 α = 0.5
4 α = 0.8
β = 0.5
β = 1.5
β = 5

Figure 3: Transitivity Coefficient. Expected C vs. expected number of clusters for different αs and βs (`continuous line`). Expected value of C under a random graph or single component model with the same $\bar{\theta}$ (`dotted line`). Numbers (1:4) represent increasing values of $\alpha = \{0, 0.2, 0.5, 0.8\}$, colors represent $\beta = 0.5$ (`blue`), 1.5 (`red`), 5 (`green`).

From now on, dyads and networks are assumed to be conditionally independent, so that

$$y_{i,i',j} \sim_{iid} \psi_j(y_{i,i',j}|\theta_{\gamma_{j,i},\gamma_{j,i'},j}), \tag{7}$$

where ψ_j is a parametric distribution associated with network j, $\theta_{k,l,j}$ is the parameter that controls the rate of interaction among factions k and l in network j, and $\gamma_{j,i}$ is the faction membership indicator for actor i in network j.

In order to identify groups of networks with similar community structures, we introduce a set of indicators $\zeta_1, \ldots, \zeta_J$ associated with each of the networks, such that $\zeta_j = \zeta_{j'}$ if and only if $\gamma_{j,i} = \boldsymbol{\gamma}_{j',i}$ for all $i = 1, \ldots, I$. Hence, inferences on $\zeta_1, \ldots, \zeta_J$ allow us to asses how similar the community

structures are across different networks. A joint prior for $\boldsymbol{\gamma}_1, \ldots, \boldsymbol{\gamma}_J$ is then obtained by setting

$$\zeta_j | \boldsymbol{\varpi} \sim_{iid} \sum_{k=1}^{\infty} \varpi_k \delta_k, \qquad \varpi_k = v_k \prod_{s<k} (1 - v_s), \qquad v_k \sim_{iid} \mathsf{Beta}(1 - \alpha_1, \beta_1 + \alpha_1 k) \tag{8}$$

and letting $\boldsymbol{\gamma}_j = \boldsymbol{\xi}_{\zeta_j}$, where vectors $\boldsymbol{\xi}_1, \boldsymbol{\xi}_2, \ldots$ (which encode the unique community structures associated with each group of networks) are independently sampled according to

$$\xi_{k,i} | \mathbf{w} \sim_{iid} \sum_{l=1}^{\infty} w_{k,l} \delta_l, \qquad w_{k,l} = u_{k,l} \prod_{s<l} (1 - u_{k,s}), \qquad u_{k,l} \sim_{iid} \mathsf{Beta}(1 - \alpha_2, \beta_2 + \alpha_2 l). \tag{9}$$

The conditional independence assumptions implicit in (8) and (9) ensure exchangeability across the networks, i.e., that the model is invariant to the order in which the networks are included in the model.

Equation (9) implies a joint distribution for $(\xi_{k,1}, \ldots, \xi_{k,J})$ that can be described by a generalized Chinese restaurant process (Pitman, 1995) that sets $\xi_{k,1} = 1$ and

$$\xi_{k,i} | \xi_{k,i-1}, \ldots, \xi_1 \sim \sum_{l=1}^{L^{i-1}} \frac{n_{k,l}^{i-1} - \alpha_2}{\beta_2 + i - 1} \delta_l + \frac{\beta_2 + \alpha_2 L^{i-1}}{\beta_2 + i - 1} \delta_{L^{i-1}+1}, \qquad 2 \leq i \leq I, \tag{10}$$

where $n_{k,l}^{i-1} = \sum_{j=1}^{i-1} \mathbf{1}_{(\xi_{k,j}=l)}$ is the number of members of faction l among the first $i-1$ subjects in the k-th cluster of networks, and $L^{i-1} = \max_{j<i}\{\xi_{k,j}\}$ is the number of factions that are represented among the first $i-1$ subjects. Note that taking $\alpha_2 = 0$ takes us back to the specification in (2).

Similar comments apply to the implied joint distribution on the cluster indicators $\zeta_1, \ldots, \zeta_J$. In particular, note that integrating out the prior weights $\mathbf{w}$ we obtain the prior probability that two networks have the same community structure, $\mathsf{Pr}(\zeta_j = \zeta_{j'} | \alpha_1, \beta_1, \alpha_2, \beta_2) = (1 - \alpha_1)/(\beta_1 + 1)$; taking $\beta_1 \to \infty$ implies that the networks are modeled independently and no information is borrowed, while taking $\beta_1 \to 0$ implies that all networks share the same community structure. In addition, note that

$$\begin{aligned} \mathsf{Pr}(\gamma_{j,i} = \gamma_{j,i'} | \gamma_{j',i} = \gamma_{j',i'}, \alpha_1, \beta_1, \alpha_2, \beta_2) &= \frac{1 - \alpha_1}{\beta_1 + 1} + \frac{\beta_1 + \alpha_1}{\beta_1 + 1} \frac{1 - \alpha_2}{\beta_2 + 1} \\ &\geq \frac{1 - \alpha_2}{\beta_2 + 1} = \mathsf{Pr}(\gamma_{j,i} = \gamma_{j,i'} | \alpha_2, \beta_2) \end{aligned}$$

with the equality happening only if $\beta_1 \to \infty$. Hence the prior probability that two subjects belong to the same faction under the joint model is strictly larger than that implied by independently modeling each network.

The model is completed by specifying a prior distribution on the interaction matrices $\boldsymbol{\Theta}_1, \ldots, \boldsymbol{\Theta}_J$.

As in Section 3 we let

$$p_j(\boldsymbol{\Theta}_j|\boldsymbol{\lambda}_j) = \begin{cases} \left\{\prod_{k=1}^{\infty} H_j^{\boldsymbol{\lambda}_{j,D}}(\theta_{k,k,j})\right\}\left\{\prod_{k=1}^{\infty}\prod_{l=k+1}^{\infty} H_j^{\boldsymbol{\lambda}_{j,O}}(\theta_{k,l,j})\delta_{\theta_{k,l,j}}(\theta_{l,k,j})\right\} & \mathbf{Y}_j \text{ undirected} \\ \left\{\prod_{k=1}^{\infty} H_j^{\boldsymbol{\lambda}_{j,D}}(\theta_{k,k,j})\right\}\left\{\prod_{k=1}^{\infty}\prod_{l=k+1}^{\infty} H_j^{\boldsymbol{\lambda}_{j,O}}(\theta_{k,l,j})H_j^{\boldsymbol{\lambda}_{j,O}}(\theta_{l,k,j})\right\} & \mathbf{Y}_j \text{ directed} \end{cases}, \tag{11}$$

where $\boldsymbol{\lambda}_j = (\boldsymbol{\lambda}_{j,D}, \boldsymbol{\lambda}_{j,O})$.

4.1 Markov chain Monte Carlo inference

The posterior distribution implied by the model described in Section 4 is

$$\begin{aligned} p(\{\boldsymbol{\Theta}_j\}, \{\boldsymbol{\xi}_k\}, \boldsymbol{\zeta}, \{\boldsymbol{\lambda}_j\}, \alpha_1, \beta_1, \boldsymbol{\alpha}_2, \boldsymbol{\beta}_2|\{\mathbf{Y}_j\}) \propto & \left\{\prod_{j=1}^{J}\prod_{i=1}^{I}\prod_{i'=1, i\neq i'}^{I} \psi_j(y_{i,i',j}|\theta_{\xi_{\zeta_j,i},\xi_{\zeta_j,i'},j})\right\} \\ & \left\{\prod_{j=1}^{J} p_j(\boldsymbol{\Theta}_j|\boldsymbol{\lambda}_j)\right\}\left\{\prod_{k=1}^{\max\{\zeta_j\}} p(\boldsymbol{\xi}_k|\alpha_{2,k}, \beta_{2,k})\right\} p(\boldsymbol{\zeta}|\alpha_1, \beta_1)\left\{\prod_{j=1}^{J} p(\boldsymbol{\lambda}_j)\right\} \\ & p(\alpha_1, \beta_1)\left\{\prod_{k=1}^{\max\{\zeta_j\}} p(\alpha_{2,k}, \beta_{2,k})\right\} \end{aligned} \tag{12}$$

where

$$p(\boldsymbol{\zeta}|\alpha_1, \beta_1) = \frac{\Gamma(\beta_1+1)}{(\beta_1+K\alpha_1)\Gamma(\beta_1+J)}\prod_{k=1}^{K}(\beta_1+k\alpha_1)\frac{\Gamma(n_k-\alpha_1)}{\Gamma(1-\alpha_1)}, \quad K = \max_{j\leq J}\{\zeta_j\}, \quad n_k = \sum_{j=1}^{J}\mathbf{1}_{(\zeta_j=k)}$$

and

$$p(\boldsymbol{\xi}_k|\alpha_{2,k}, \beta_{2,k}) = \frac{\Gamma(\beta_{2,k}+1)}{(\beta_{2,k}+L_k\alpha_{2,k})\Gamma(\beta_{2,k}+I)}\prod_{l=1}^{L_k}(\beta_{2,k}+l\alpha_{2,k})\frac{\Gamma(m_{k,l}-\alpha_{2,k})}{\Gamma(1-\alpha_{2,k})} \tag{13}$$

with $L_k = \max_{i\leq I}\{\xi_{k,i}\}$, $m_{k,l} = \sum_{i=1}^{I}\mathbf{1}_{(\xi_{k,i}=l)}$. This posterior distribution is computationally intractable, even if the baseline measure $H_j^{\boldsymbol{\lambda}}$ is chosen to be conjugate to the kernel ψ_j. Indeed, the number of possible groups of networks and the number of factions within each group of networks grows exponentially fast with J and I, making it impossible to explicitly enumerate all possible models.

To overcome this difficulty we develop a Markov chain Monte Carlo (MCMC) sampler to jointly explore the posterior distribution in (12). We exploit the conjugacy of ψ_j and H_j and factorize the

posterior distribution as

$$p(\{\boldsymbol{\Theta}_j\}, \{\boldsymbol{\xi}_k\}, \boldsymbol{\zeta}, \{\boldsymbol{\lambda}_j\}, \alpha_1, \beta_1, \boldsymbol{\alpha}_2, \boldsymbol{\beta}_2 | \{\mathbf{Y}_j\}) = \\ p(\{\boldsymbol{\Theta}_j\} | \{\boldsymbol{\xi}_k\}, \boldsymbol{\zeta}, \{\boldsymbol{\lambda}_j\}, \{\mathbf{Y}_j\}) p(\{\boldsymbol{\xi}_k\}, \boldsymbol{\zeta}, \{\boldsymbol{\lambda}_j\}, \alpha_1, \beta_1, \boldsymbol{\alpha}_2, \boldsymbol{\beta}_2 | \{\mathbf{Y}_j\}). \quad (14)$$

Sampling from $p(\{\boldsymbol{\Theta}_j\} | \{\boldsymbol{\xi}_k\}, \boldsymbol{\zeta}, \{\boldsymbol{\lambda}_j\}, \alpha_1, \beta_1, \boldsymbol{\alpha}_2, \boldsymbol{\beta}_2, \{\mathbf{Y}_j\})$ is straightforward under conjugacy. To sample from the marginal posterior $p(\{\boldsymbol{\xi}_k\}, \boldsymbol{\zeta}, \{\boldsymbol{\lambda}_j\}, \alpha_1, \beta_1, \boldsymbol{\alpha}_2, \boldsymbol{\beta}_2 | \{\mathbf{Y}_j\})$ we iteratively sample from the following six sets of full conditional distributions

1. $p(\boldsymbol{\zeta} | \{\boldsymbol{\xi}_k\}, \{\boldsymbol{\lambda}_j\}, \alpha_1, \beta_1, \boldsymbol{\alpha}_2, \boldsymbol{\beta}_2, \{\mathbf{Y}_j\})$.
2. $p(\{\boldsymbol{\xi}_k\}, \boldsymbol{\zeta} | \{\boldsymbol{\lambda}_j\}, \alpha_1, \beta_1, \boldsymbol{\alpha}_2, \boldsymbol{\beta}_2, \{\mathbf{Y}_j\})$.
3. $p(\boldsymbol{\xi}_k | \boldsymbol{\zeta}, \{\boldsymbol{\lambda}_j\}, \alpha_1, \beta_1, \boldsymbol{\alpha}_2, \boldsymbol{\beta}_2, \{\mathbf{Y}_j\})$ for $k = 1, \ldots, \max_{j \leq J}\{\zeta_j\}$.
4. $p(\{\boldsymbol{\lambda}_j\} | \{\boldsymbol{\xi}_k\}, \boldsymbol{\zeta}, \{\mathbf{Y}_j\})$ for $j = 1, \ldots, J$.
5. $p(\alpha_1, \beta_1 | \boldsymbol{\zeta})$.
6. $p(\{\alpha_{2,k}, \beta_{2,k}\} | \boldsymbol{\xi}_k)$ for $k = 1, \ldots, \max_{j \leq J}\{\zeta_j\}$.

To sample from $p(\{\boldsymbol{\xi}_k\}, \boldsymbol{\zeta} | \{\boldsymbol{\lambda}_j\}, \alpha_1, \beta_1, \boldsymbol{\alpha}_2, \boldsymbol{\beta}_2, \{\mathbf{Y}_j\})$ we develop a split-merge algorithm that combines ideas from Dahl (2003) and Jain & Neal (2004). More specifically, at each iteration of the MCMC we randomly select two networks. If they currently belong to two separate clusters we propose to merge them into a single cluster. On the other hand, if they belong to the same cluster we propose to split it into two different clusters. The faction structure within each cluster of networks is proposed by sequentially allocating actors to factions, in the spirit Dahl (2003). A similar approach is employed to sample from $p(\boldsymbol{\xi}_k | \boldsymbol{\zeta}, \{\boldsymbol{\lambda}_j\}, \alpha_1, \beta_1, \boldsymbol{\alpha}_2, \boldsymbol{\beta}_2, \{\mathbf{Y}_j\})$. These long-range moves are combined with more traditional short-range moves that individually update each component of $\boldsymbol{\zeta}$ and $\boldsymbol{\xi}_k$ given the rest of the components. Details on the algorithm can be seen in Appendix F.

The posterior distributions associated with $\boldsymbol{\zeta}$ and $\{\boldsymbol{\xi}_k\}$ can be summarized through posterior pairwise incidence matrices. For example, the pairwise incidence matrix associated with $\boldsymbol{\zeta}$, $\mathbf{D}_\zeta$ is an $J \times J$ matrix such that $[\mathbf{D}_\zeta]_{j,j'} = \mathsf{Pr}(\zeta_j = \zeta_{j'} | \mathbf{Y}_1, \ldots, \mathbf{Y}_J)$. To obtain point estimates of the partition structure we follow Lau & Green (2007) and take a decision theoretic approach. For example, a point estimator $\tilde{\boldsymbol{\zeta}}$ for $\boldsymbol{\zeta}$ is obtained by minimizing the expected loss function

$$L(\tilde{\boldsymbol{\zeta}}) = \mathsf{E}\left\{ \sum_{j=1}^{J} \sum_{j'=j+1}^{J} \left[a\mathbf{1}_{(\zeta_j = \boldsymbol{\zeta}_{j'}, \tilde{\zeta}_j \neq \tilde{\boldsymbol{\zeta}}_{j'})} + b\mathbf{1}_{(\zeta_j \neq \boldsymbol{\zeta}_{j'}, \tilde{\zeta}_j = \tilde{\boldsymbol{\zeta}}_{j'})} \right] \mid \mathbf{Y}_1, \ldots, \mathbf{Y}_J \right\} \quad (15)$$

Minimizing (15) is equivalent to maximizing

$$U(\tilde{\boldsymbol{\zeta}}) = \sum_{j=1}^{J} \sum_{j'=j+1}^{J} \mathbf{1}_{(\tilde{\zeta}_j = \tilde{\boldsymbol{\zeta}}_{j'})} \left\{ \mathsf{Pr}(\zeta_j = \zeta_{j'} | \mathbf{Y}_1, \ldots, \mathbf{Y}_J) - \frac{b}{a+b} \right\}$$

The constants a and b represent the costs of misclassification errors; setting $b = 0$ leads to a point estimate that includes all networks into a single partition, while $a = 0$ leads to a point estimate the places each network into an individual partition.

5 Illustrations

5.1 Simulation study

A simulation study was conducted to compare the results between modeling a collection of networks individually applying model (1) or using the model proposed in Section 4. Seven networks with 21 actors each were simulated. The first four are binary and undirected, the fifth is Poisson undirected, and the last two are binary and directed (this setting is similar to the real data we consider in Section 5.3). The simulated networks group as follow (1,6), (2,5,7), (3), (4).

Figure 4 presents the results in terms of the estimated posterior probability that two actors belong to the same group. Comparing in particular the first networks of classes 1 and 2, we can observe that those networks borrow strength from the other networks on the group to get a better estimation of the factions when the multiple network model is used. Moreover, by looking to the estimation on left-hand side of figure 4 alone, it would have been difficult to identify the similitude of the networks 1 and 2 to the rest in classes 1 and 2, respectively. For classes 4 and 5, there is not sufficient information on the data to detect the true factions, notice that in this case both modeling techniques are equivalent, and considering MCMC error, they lead to similar results as expected.

This experiment provides evidence that in the presence of similar factions for two or more networks, the estimation benefits from modeling multiple networks simultaneously. Moreover, the model does not force artificial groupings of networks and eventually leads to equivalent results when networks need to be modeled independently.

5.2 Modeling cognitive social structures: The Krackhardt dataset

Krackhardt (1987) presents information about the relationship among management personnel of a small manufacturing company producing high-tech machinery in the west coast of the U.S. At the

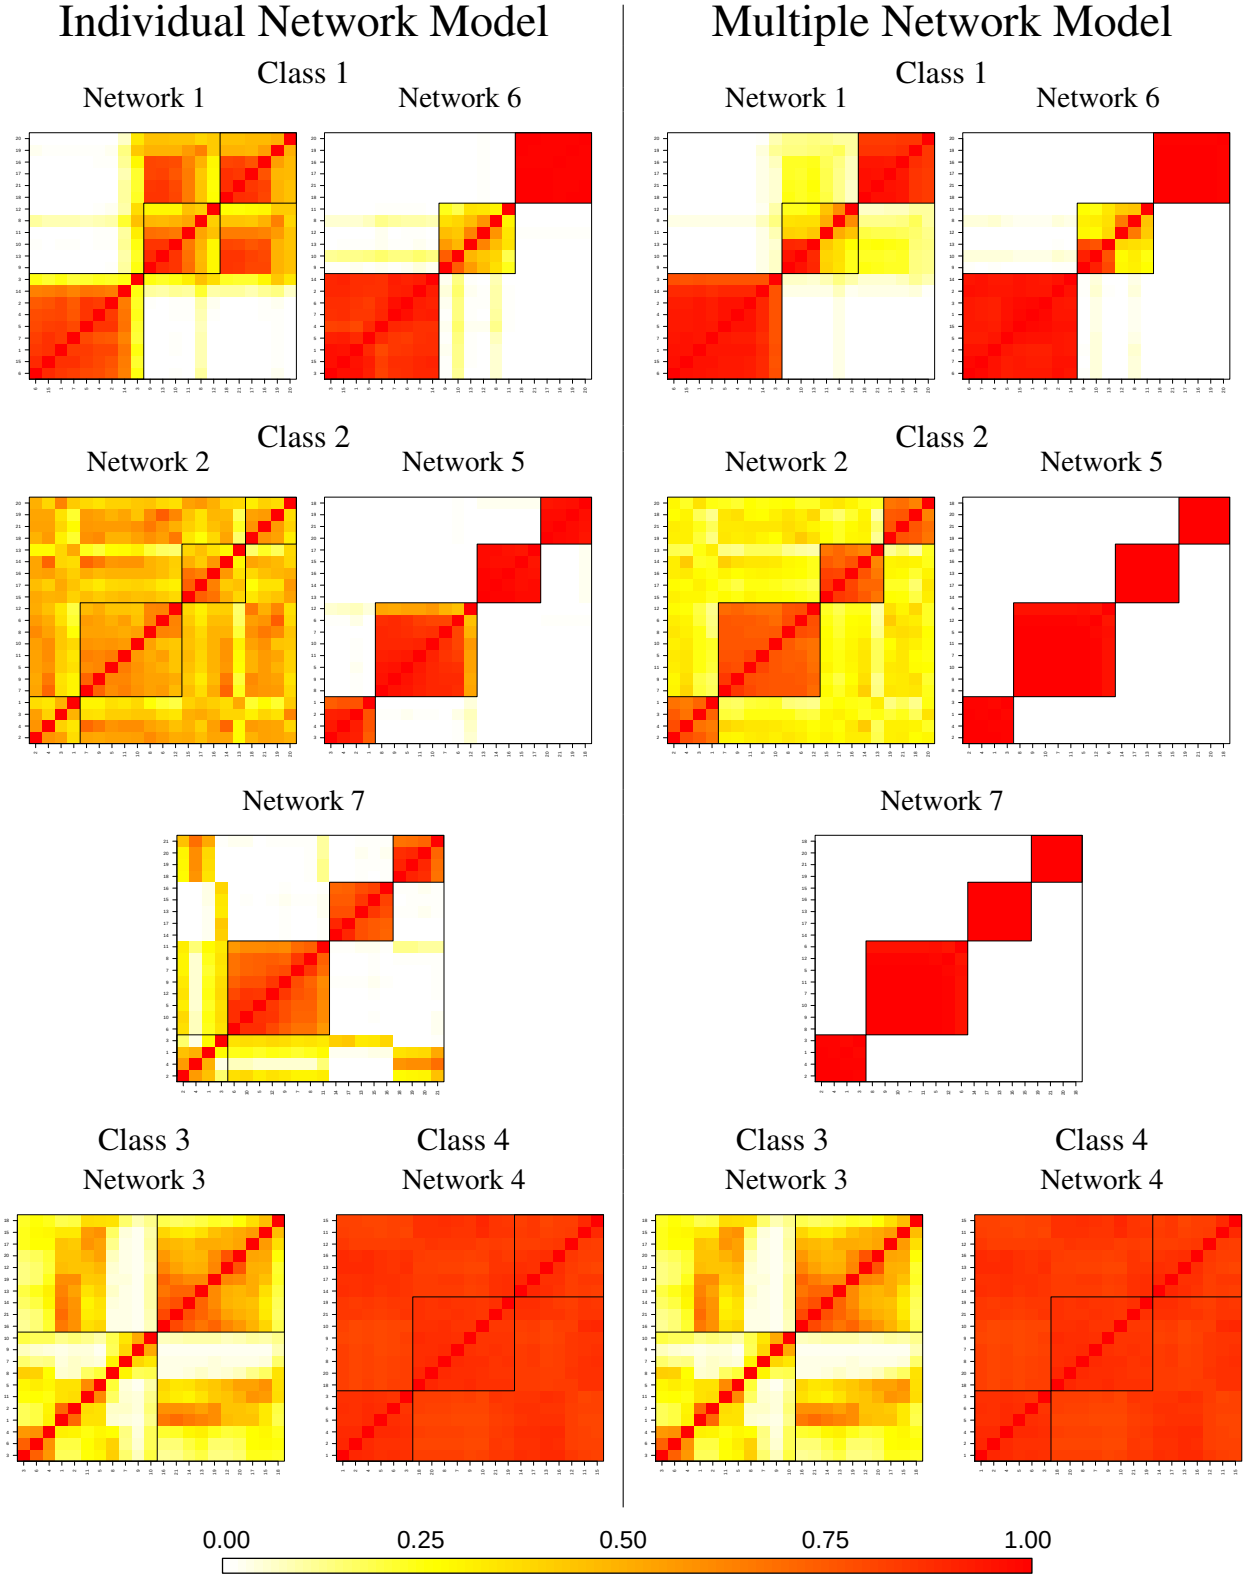

Figure 4: Simulation Results. Estimated posterior probability that two actors belong to the same group for each network. The squares mark the real groups used to create the data.

Table 1: Krackhardt data. Manager's attributes by class

Class	Actor	Age	Tenure	Department	Level
1	4	33	7.5	4	3
	18	33	9.1	3	2
	21	36	12.5	1	2
	16	27	4.7	4	3
	2	42	19.6	4	2
	8	34	11.3	1	3
	12	34	8.9	1	3
	1	33	9.3	4	3
	15	40	8.4	2	3
2	3	40	12.8	2	3
	9	62	5.4	2	3
3	5	32	3.3	2	3
	19	32	4.8	2	3
	14	43	10.4	2	2
4	6	59	28.0	1	3
5	7	55	30.0	0	1
6	10	37	9.3	3	3
7	11	46	27.0	3	3
8	13	48	0.3	2	3
9	17	30	12.4	1	3
10	20	38	11.7	2	3

time, it had about one hundred employees, including 21 managers (CEO, 4 Vice-presidents and 16 managers) that are the set of actors of this example. We focus on friendship relationships; each manager was asked not only about their connections ("Who are your friends?"), but also about their perception of other manager's connections ("Who is a friend of ... ?"), resulting in 21 networks, one for each actor's perception of the friendship network.

The multiple networks model presented in Section 4 groups networks accordingly with how actors form factions in each network. Three classes were identified with 9, 2 and 3 networks, while the other 7 networks were left ungrouped. Many of these singleton clusters arise because the individuals see themselves as friends of everybody else, a vision that is not shared by the other members of the network. Table 1 presents additional information on the managers listed accordingly with the classes obtained by the proposed technique. The four attributes for the managers obtained from table B.4 in Wasserman & Faust (1994) are age, tenure (time employed by the company, in years), department and level in the hierarchy (1: CEO, 2: vice-presidents and 3: managers). Note that all the managers in department 4 belong to class one, while most of department 2 subjects are in classes two and three. Class 3 is formed by department 2's manager and its two newest employees. In general, newer

employees appear to share the perception of others, either with more tenure or with more authority. The only exception, subject 13 may be too new in the company to have been influenced.

Figure 5 shows the resulting community structures for one representative actor from each class. Although all plots have the same ordering and inner squares have been added to help with comparisons, neither of them are the result of a formal estimation procedure. Class one actors described the network having several connections, the actors tend to be grouped in factions of friendship groups (more connected within groups than between). Class two subjects indicate very few friendship connections in their networks. Class three shows more connections in their networks and a tendency of forming groups according to friendship patterns. The ungrouped managers showed diverse oddities that made their networks different than the rest, subject 17 is an example of one of the most common of this differences. Note that, as we highlighted before, he marked himself as being friend of almost all of the managers.

In addition to comparing actors' perceptions of community structures, we studied the assortativity for each of the 21 networks. As we discussed in Section 3, a simple summary of this property is the assortativity index Υ defined in section 3.2. Figure shows boxplots of the posterior distributions of Υ for each of the 21 networks, along with a sample from the prior distribution (which was identical for all networks). Evidently all actors describe the network as assortative, i.e. interactions among actors in the same faction tend to be more common than interactions among actors in different factions. This is exactly what we expected to see for this type of relationship.

5.3 The Wiring dataset

Finally, we applied our model to a dataset consisting of the interactions among 14 Western Electric (Hawthorne Plant) employees working on the bank wiring room (Roethlisberger & Dickson, 1939). The employees worked in a single room and six relationships were recorded: friendship (binary, undirected), participation in horseplay (binary, undirected), helping others with work (binary, directed), antagonistic behavior (binary, undirected), arguments about open windows (binary, directed), and number of times workers traded job assignments (integer valued, directed).

The posterior distribution on ζ strongly suggests that the community structure associated with each network is different from the rest. Indeed, the highest posterior probability that two relationships have the same structure is 0.23 (for friendship and participation in horseplay). Figure 7 shows the estimated mean posterior probabilities/intensities and tentative groupings for each network. Most relationships form clear factions with the exception of Help, where the groups are not as delimited. Note that by clustering networks according to their underlying community structures we are able to

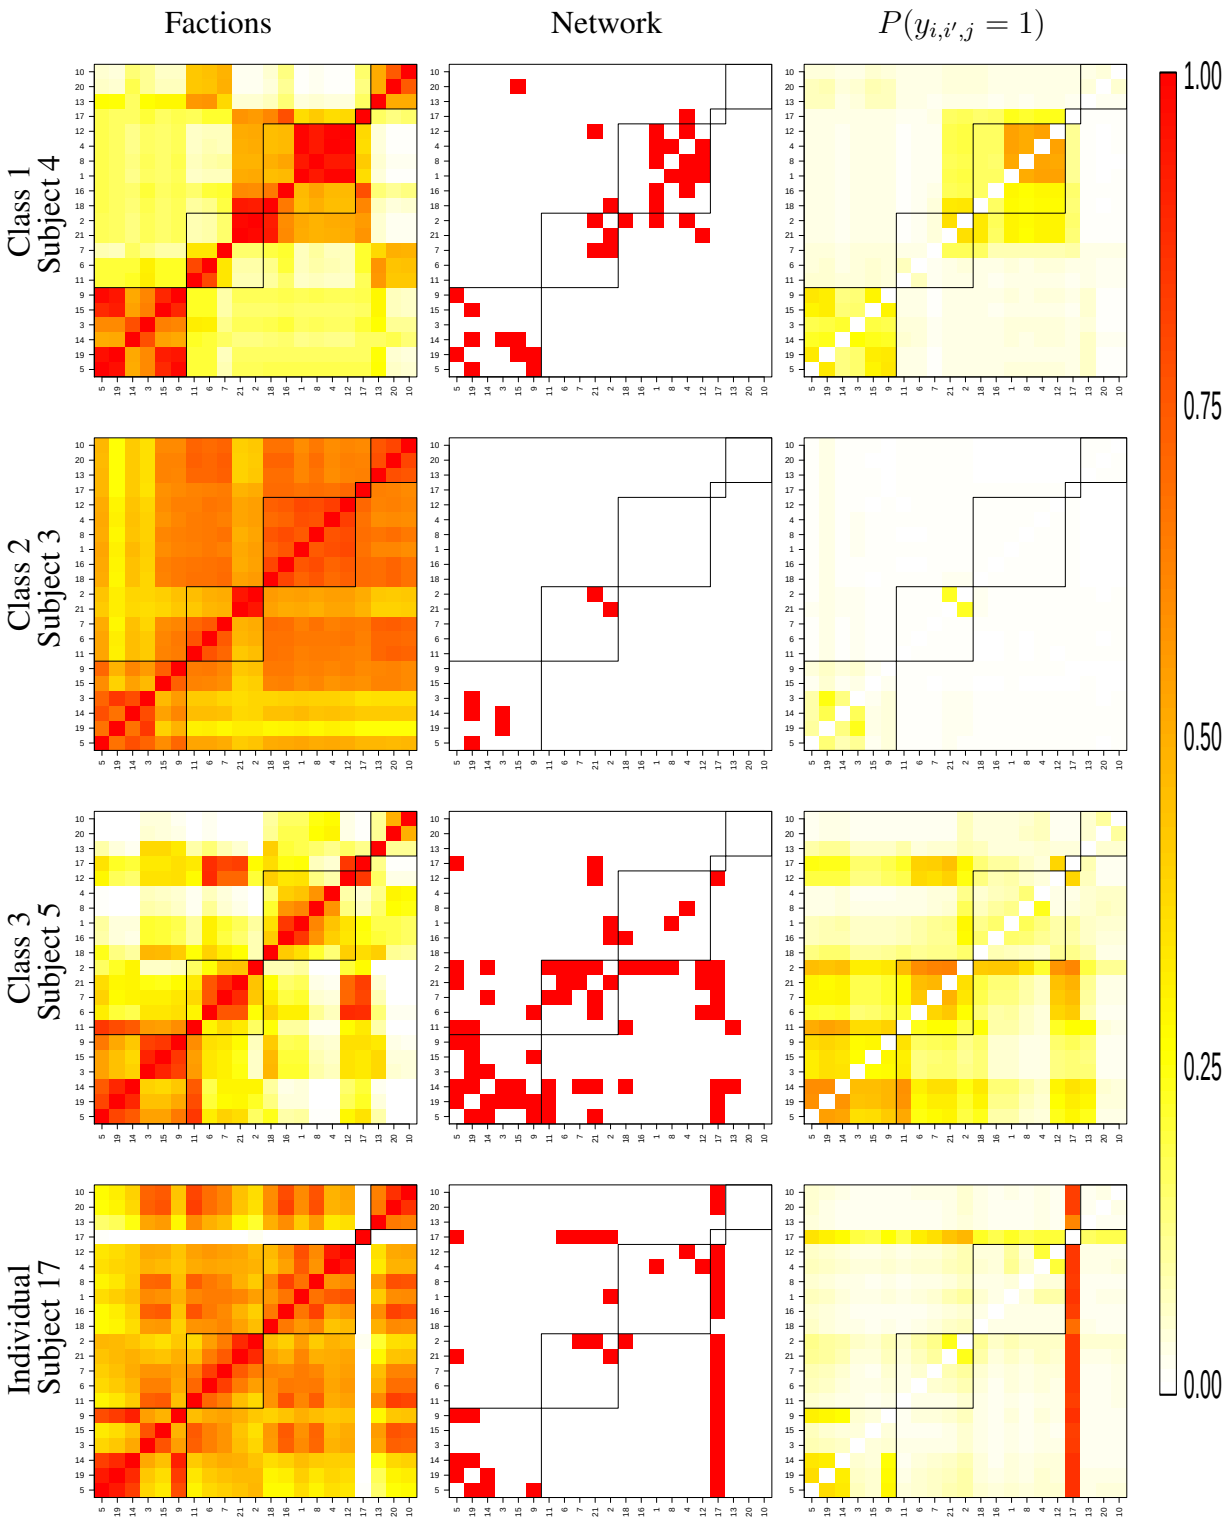

Figure 5: Krackhardt data. Individual community structures (estimated probability of being in the same community), friendship network, and estimated probability of connection between actors.

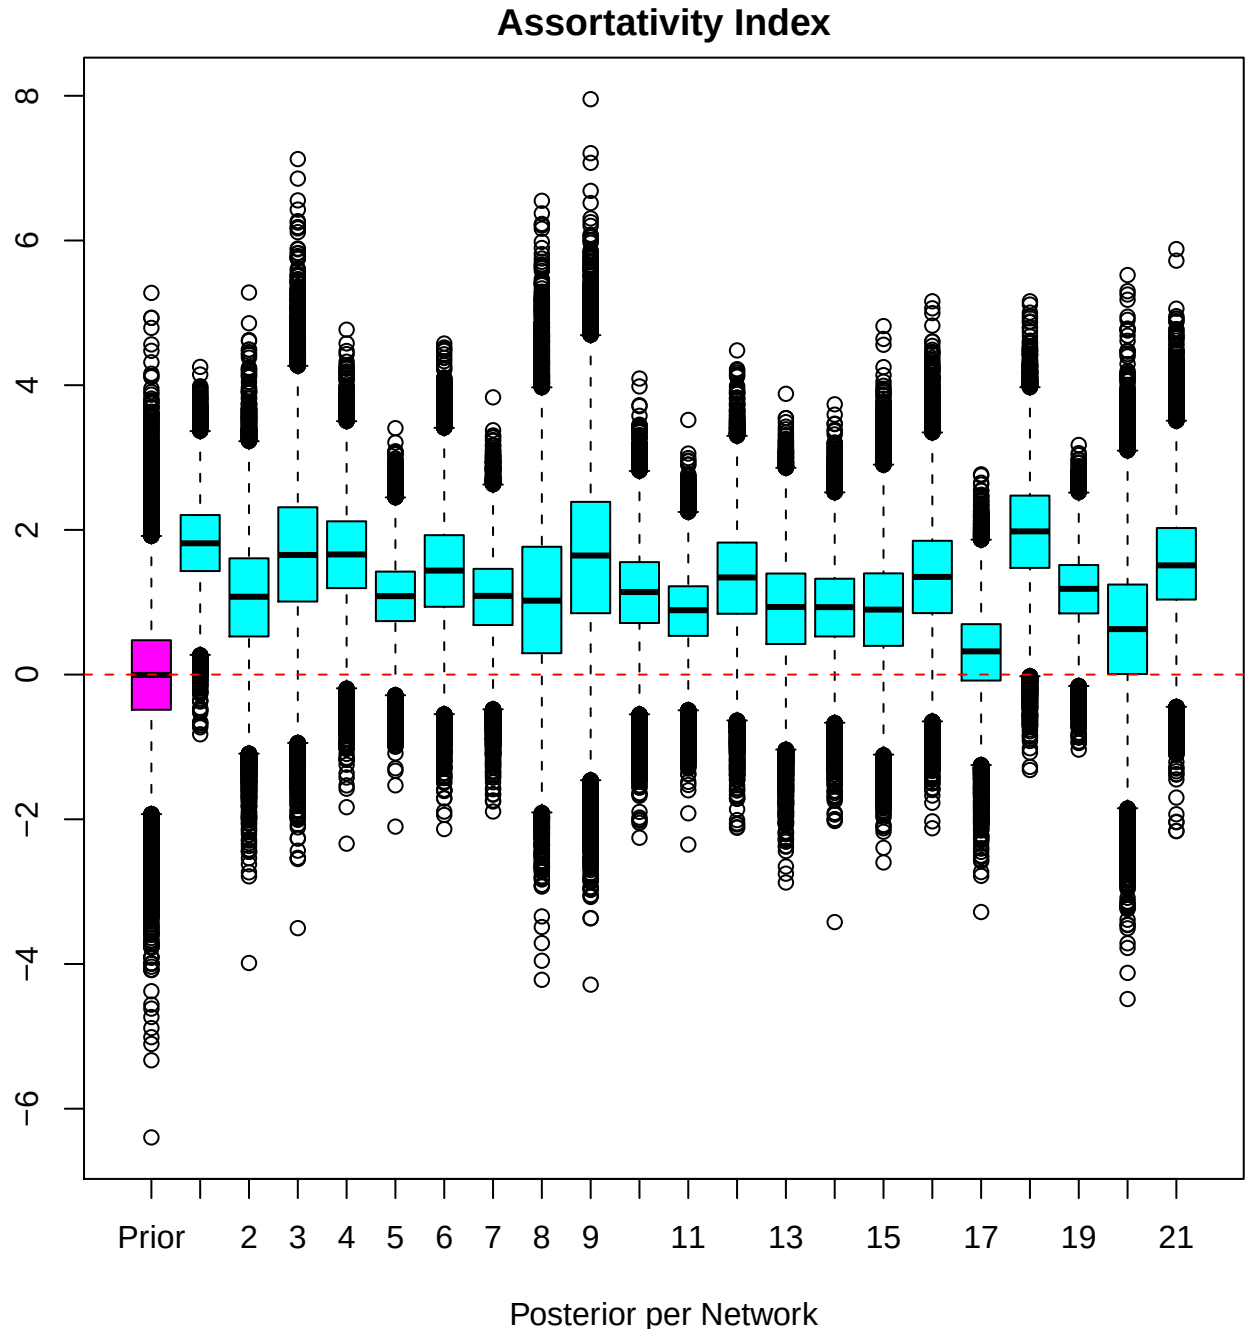

Figure 6: Krackhardt data. Prior (`left`) and Posterior distribution of the assortativity index per network

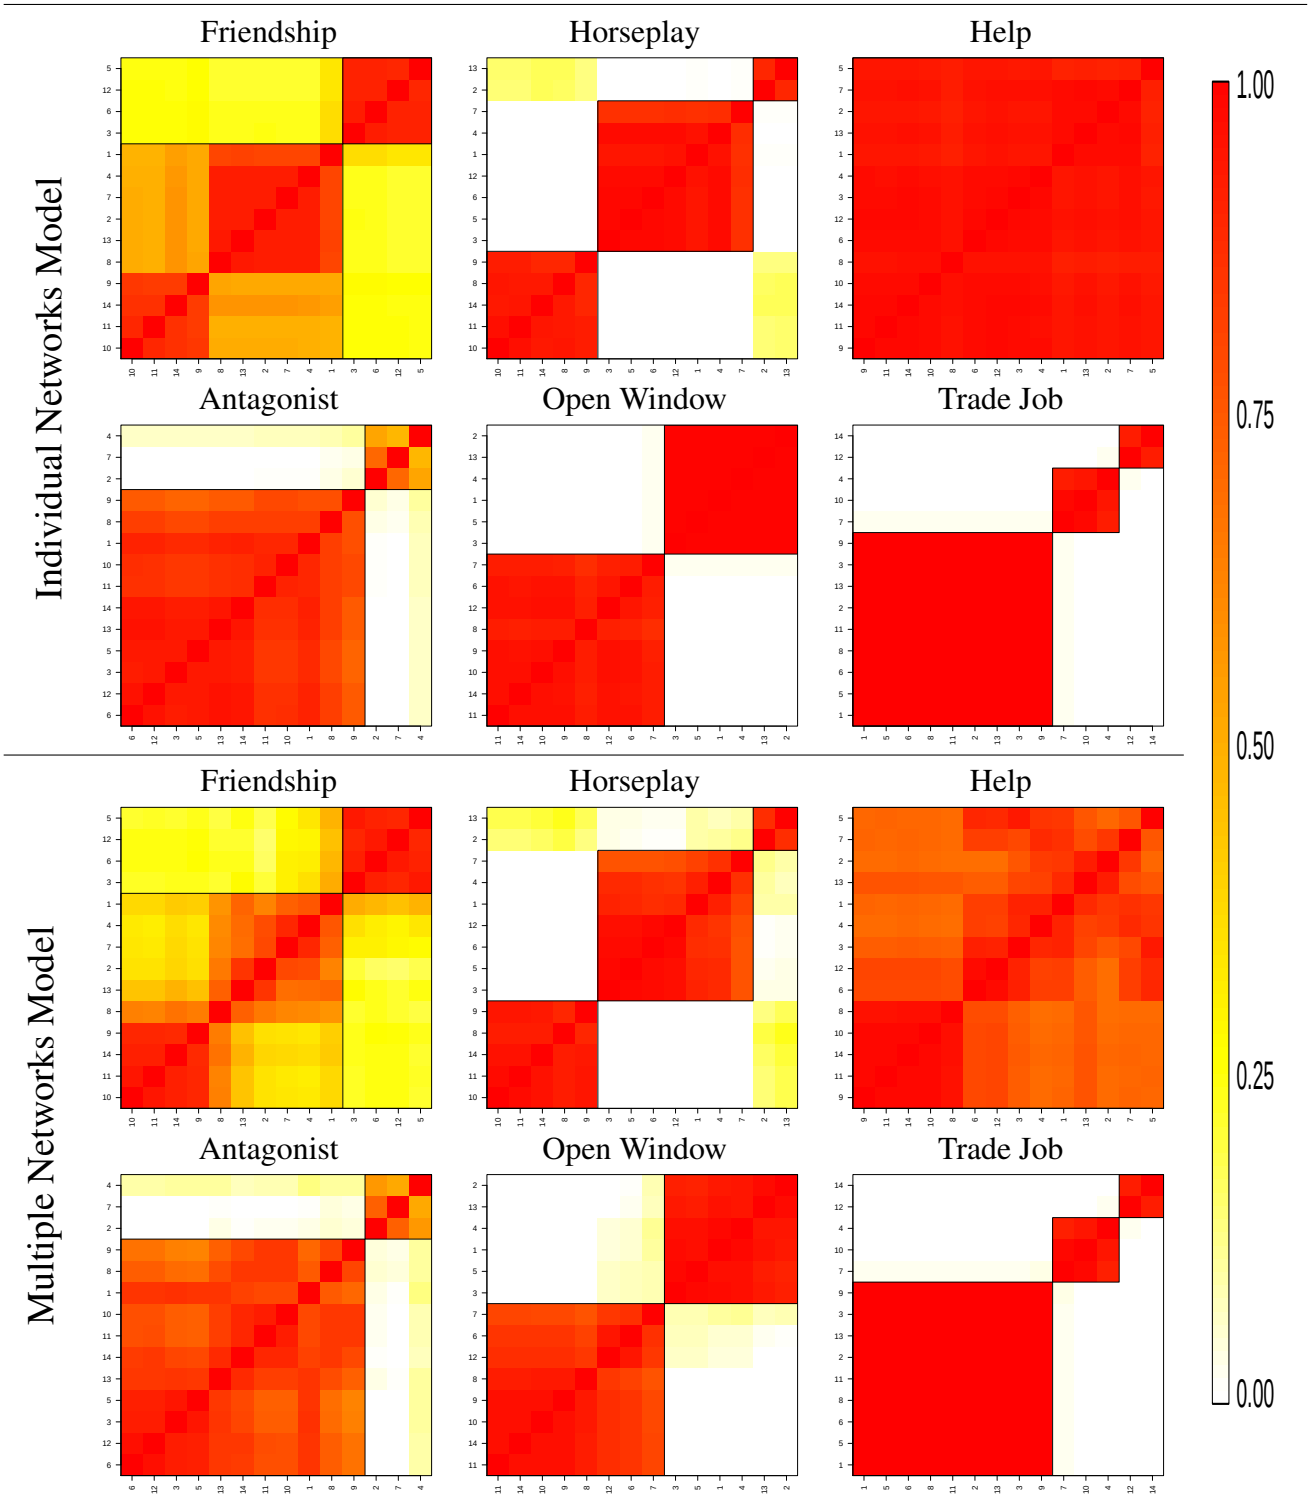

Figure 7: Wiring dataset. Estimated posterior probability that two actors belong to the same group for each network.

simultaneously compare directed, undirected, binary and valued networks.

6 Conclusions

This manuscript makes two contributions to the literature on network analysis. Firstly, we have described a general class of priors for blockmodels that allow us to incorporate available prior information about the structure of the network and construct model-based summary statistics for properties such as assortativity. Secondly, we have constructed a joint model for multiple networks that allows us to compare their underlying social structures, even across networks of different types.

With regard to the first contribution, it is worthwhile noting that further generalizations of the model are possible. For example, by constructing a prior for $\mathbf{\Theta}$ where the entries for $\theta_{k,l}$ adn $\theta_{l,k}$ are identically distributed but dependent, the model can accommodate reflexivity as well more flexible specifications for the correlation among the in-degree and out-degree of directed networks. With regard to the second contribution, the results from our analyses suggests that, for some applications, assuming that the community structure associated with various networks is the same might be too strong. A model with softer constrains is presented in Rodriguez et al. (2011), however, such model is not helpful when dealing with networks of different types. In this regard, an interesting open problem is constructing joint distributions for multiple partitions of a given set of items; such models can in turn be used to create more flexible models for collections of networks.

A Proof of lemma 1

We focus on the case of directed networks, the proof for undirected networks follows along the same lines. For any fixed K, consider the $K \times K$ matrix $\mathbf{\Theta} = [\theta_{k,k}]$ with joint distribution

$$p(\mathbf{\Theta}) = \left\{ \prod_{k=1}^{K} H^{\boldsymbol{\lambda}_D}(\theta_{k,k}) \right\} \left\{ \prod_{k=1}^{K} \prod_{l=k+1}^{K} H^{\boldsymbol{\lambda}_O}(\theta_{k,l}) H^{\boldsymbol{\lambda}_O}(\theta_{l,k}) \right\}.$$

To show that $p(\mathbf{\Theta})$ is invariant to any permutation, it is enough to show that it is invariant to permutations that only exchange any two indexes i and j, i.e., it is invariant to the permutation $\sigma(i) = j$, $\sigma(j) = i$, and $\sigma(k) = k$ for any other k different from i and j. This is because any general permutation can be written as a composition of these simple permutations. Let $\mathbf{\Theta}^*$ be the permuted version of

$\boldsymbol{\Theta}$. Now

$$p(\boldsymbol{\Theta}^*) = \left\{ \prod_{k=1}^{K} H^{\boldsymbol{\lambda}_D}(\theta^*_{k,k}) \right\} \left\{ \prod_{k=1}^{K} \prod_{l=k+1}^{K} H^{\boldsymbol{\lambda}_O}(\theta^*_{k,l}) H^{\boldsymbol{\lambda}_O}(\theta^*_{l,k}) \right\}.$$

Note that $H^{\boldsymbol{\lambda}_D}(\theta^*_{i,i}) = H^{\boldsymbol{\lambda}_D}(\theta_{\sigma(i),\sigma(i)}) = H^{\boldsymbol{\lambda}_D}(\theta_{j,j})$ and $H^{\boldsymbol{\lambda}_D}(\theta^*_{j,j}) = H^{\boldsymbol{\lambda}_D}(\theta_{\sigma(j),\sigma(j)}) = H^{\boldsymbol{\lambda}_D}(\theta_{ii})$, while

$$H^{\boldsymbol{\lambda}_D}(\theta^*_{k,k}) = H^{\boldsymbol{\lambda}_D}(\theta_{\sigma(k),\sigma(k)}) = H^{\boldsymbol{\lambda}_D}(\theta_{k,k})$$

for any $k \neq i, j$. Hence,

$$\prod_{k=1}^{K} H^{\boldsymbol{\lambda}_D}(\theta^*_{k,k}) = \prod_{k=1}^{K} H^{\boldsymbol{\lambda}_D}(\theta_{k,k}).$$

A similar argument can be made for the off-diagonal elements. In this case, the only terms that are not obviously identical before and after the permutation are those associated with $(\theta_{i,k}, \theta_{k,i})$ for $k \neq i, j$, those associated with $(\theta_{j,k}, \theta_{k,j})$ for $k \neq i, j$, and the term associated with $(\theta_{i,j}, \theta_{i,j})$. Since the pairs are independent, it is clear that

$$\begin{aligned}
\prod_{k\neq i,j} H^{\boldsymbol{\lambda}_O}(\theta^*_{i,k}) H^{\boldsymbol{\lambda}_O}(\theta^*_{k,i}) \prod_{k\neq i,j} H^{\boldsymbol{\lambda}_O}(\theta^*_{j,k}) H^{\boldsymbol{\lambda}_O}(\theta^*_{k,j}) = \\
\prod_{k\neq i,j} H^{\boldsymbol{\lambda}_O}(\theta_{\sigma(i),k}) H^{\boldsymbol{\lambda}_O}(\theta_{k,\sigma(i)}) \prod_{k\neq i,j} H^{\boldsymbol{\lambda}_O}(\theta_{\sigma(j),k}) H^{\boldsymbol{\lambda}_O}(\theta_{k,\sigma(j)}) = \\
\prod_{k\neq i,j} H^{\boldsymbol{\lambda}_O}(\theta_{j,k}) H^{\boldsymbol{\lambda}_O}(\theta_{k,j}) \prod_{k\neq i,j} H^{\boldsymbol{\lambda}_O}(\theta_{i,k}) H^{\boldsymbol{\lambda}_O}(\theta_{k,i}).
\end{aligned}$$

Finally, $H^{\boldsymbol{\lambda}_O}(\theta^*_{i,j}) H^{\boldsymbol{\lambda}_O}(\theta^*_{j,i}) = H^{\boldsymbol{\lambda}_O}(\theta_{\sigma(i),\sigma(j)}) H^{\boldsymbol{\lambda}_O}(\theta_{\sigma(j),\sigma(i)}) = H^{\boldsymbol{\lambda}_O}(\theta_{i,j}) H^{\boldsymbol{\lambda}_O}(\theta_{j,i})$, which completes the proof.

B Proof of lemma 2

A simple conditioning argument implies that

$$\bar{\theta} = \mathsf{E}\left\{\mathsf{E}\left(y_{ij}|\xi_i, \xi_j, \boldsymbol{\Theta}\right)\right\} = \mathsf{E}_{H^{\boldsymbol{\lambda}_D}}\{\theta_{k,k}\}\mathsf{Pr}(\xi_i = \xi_j) + \mathsf{E}_{H^{\boldsymbol{\lambda}_O}}\{\theta_{k,l}\}\mathsf{Pr}(\xi_i \neq \xi_j)$$

Now, because of the exchangeability of the observations we have $\mathsf{Pr}(\xi_i = \xi_j) = \mathsf{Pr}(\xi_1 = \xi_2) = \sum_{k=1}^{\infty} \mathsf{E}\{\omega_k^2\} = (1-\alpha)/(\beta+1)$.

C Proof of lemma 3

Again, we apply conditioning arguments. For $\bar{\rho}$,

$$\begin{aligned}\bar{\rho} = \mathsf{E}\{D_i|\boldsymbol{\lambda},\boldsymbol{\varsigma}\} &= \mathsf{E}\left\{\sum_{j\neq i} y_{ij}\middle|\boldsymbol{\lambda},\boldsymbol{\varsigma}\right\} \\ &= \sum_{j\neq i}\mathsf{E}\left\{\mathsf{E}(y_{ij}|\xi_i,\xi_j,\boldsymbol{\Theta})\right\} \\ &= [I-1]\left[\mathsf{Pr}(\xi_i=\xi_j)\mathsf{E}_{H^{\boldsymbol{\lambda}_D}}\{\theta_{k,k}\}+\mathsf{Pr}(\xi_i\neq\xi_j)\mathsf{E}_{H^{\boldsymbol{\lambda}_O}}\{\theta_{k,l}\}\right],\end{aligned}$$

where the last equality follows from lemma 2.

There are a few alternative ways to obtain this result that might be of interest. For example, let D be the number of links in the network $\mathbf{Y}$ composed of I subjects, L be the number of factions in the network and m_l for $l=1,\ldots,L$ the size of the l-th community, then

$$\mathsf{E}\{D|\boldsymbol{\lambda},\boldsymbol{\varsigma}\} = \mathsf{E}\left\{\mathsf{E}\left\{D|L,\{m_l\},\boldsymbol{\Theta}\right\}\right\} = \mathsf{E}\left\{\sum_{l=1}^{L}\sum_{k\neq l,k=1}^{L} m_l m_k\theta_{k,l}+\sum_{l=1}^{L} m_l(m_l-1)\theta_{ll}\right\}.$$

The number of factions and the sizes of the factions are determined exclusively by the weights $\{w_k\}$, hence they are independent of θ's and

$$\begin{aligned}\mathsf{E}\{D|\boldsymbol{\lambda},\boldsymbol{\varsigma}\} &= \sum_{l=1}^{L}\sum_{k\neq l,k=1}^{L}\mathsf{E}_{H^{\boldsymbol{\lambda}_O}}\{\theta_{k,l}\}\mathsf{E}\left(m_l m_k\right)+\sum_{l=1}^{L}\mathsf{E}_{H^{\boldsymbol{\lambda}_D}}\{\theta_{ll}\}\mathsf{E}\left(m_l^2-m_l\right) \\ &= \mathsf{E}_{H^{\boldsymbol{\lambda}_O}}\{\theta_{k,l}\}\mathsf{E}\left\{\sum_{l=1}^{L}\sum_{k\neq l,k=1}^{L} m_l m_k\right\}+\mathsf{E}_{H^{\boldsymbol{\lambda}_D}}\{\theta_{ll}\}\left\{\mathsf{E}\left(\sum_{l=1}^{L} m_l^2\right)-E\left(\sum_{l=1}^{L} m_l\right)\right\} \\ &= \mathsf{E}_{H^{\boldsymbol{\lambda}_O}}\{\theta_{k,l}\}\left\{I^2-\mathsf{E}\left(\sum_{l=1}^{L} m_l^2\right)\right\}+\mathsf{E}_{H^{\boldsymbol{\lambda}_D}}\{\theta_{ll}\}\left\{\mathsf{E}\left(\sum_{l=1}^{L} m_l^2\right)-I\right\}.\end{aligned}$$

In the above expression, we make repeated use of the fact that $\sum_{l=1}^{L} m_l = I$ almost surely. Hence,

$$\mathsf{E}\{D|\boldsymbol{\lambda},\boldsymbol{\varsigma}\} = I^2\mathsf{E}_{H^{\boldsymbol{\lambda}_O}}\{\theta_{k,l}\}-I\mathsf{E}_{H^{\boldsymbol{\lambda}_D}}\{\theta_{k,k}\}+\left(\mathsf{E}_{H^{\boldsymbol{\lambda}_D}}\{\theta_{k,k}\}-\mathsf{E}_{H^{\boldsymbol{\lambda}_O}}\{\theta_{k,l}\}\right)\mathsf{E}\left\{\sum_{l=1}^{L} m_l^2\right\}.$$

Finally, since observations are exchangeable the mean number of links per observation is simply

$$\bar{\rho} = \mathsf{E}\{D_i|\boldsymbol{\lambda},\boldsymbol{\varsigma}\} = \mathsf{E}\left\{\frac{D}{I}\middle|\boldsymbol{\lambda},\boldsymbol{\varsigma}\right\} = I\mathsf{E}_{H^{\boldsymbol{\lambda}_O}}\{\theta_{k,l}\}-\mathsf{E}_{H^{\boldsymbol{\lambda}_D}}\{\theta_{k,k}\}+\frac{1}{I}\left(\mathsf{E}_{H^{\boldsymbol{\lambda}_D}}\{\theta_{k,k}\}-\mathsf{E}_{H^{\boldsymbol{\lambda}_O}}\{\theta_{k,l}\}\right)\mathsf{E}\left\{\sum_{l=1}^{L} m_l^2\right\}. \quad (16)$$

An alternative way to compute $\bar{\rho}$ is working directly with D_1, without lost of generality assume that subject 1 falls into the first community,

$$\bar{\rho} = \mathsf{E}\{D_1|\boldsymbol{\lambda}, \boldsymbol{\varsigma}\} = \mathsf{E}\left\{\mathsf{E}\left(D_1|L, \{m_l\}, \boldsymbol{\Theta}\right)\right\} = \mathsf{E}\left\{(m_1 - 1)\theta_{1,1} + \sum_{l=2}^{L} m_l\theta_{1l}\right\}$$
$$= \mathsf{E}\{m_1\}\mathsf{E}_{H^{\boldsymbol{\lambda}_D}}\{\theta_{k,k}\} - \mathsf{E}_{H^{\boldsymbol{\lambda}_D}}\{\theta_{k,k}\} + \mathsf{E}\left\{\sum_{l=2}^{L} m_l\right\}\mathsf{E}_{H^{\boldsymbol{\lambda}_O}}\{\theta_{k,l}\}$$

After regrouping and using the fact that $\sum_{l=2}^{L} m_l = I - m_1$ almost surely,

$$\bar{\rho} = I\mathsf{E}_{H^{\boldsymbol{\lambda}_O}}\{\theta_{k,l}\} - \mathsf{E}_{H^{\boldsymbol{\lambda}_D}}\{\theta_{k,k}\} + \left(\mathsf{E}_{H^{\boldsymbol{\lambda}_D}}\{\theta_{k,k}\} - \mathsf{E}_{H^{\boldsymbol{\lambda}_O}}\{\theta_{k,l}\}\right)\mathsf{E}\{m_1\}. \tag{17}$$

Perhaps the simplest expression $\mathsf{E}\{D_i|\boldsymbol{\lambda}, \boldsymbol{\varsigma}\}$ was already derived in lemma 3, rearranging terms

$$\bar{\rho} = I\mathsf{E}_{H^{\boldsymbol{\lambda}_O}}\{\theta_{k,l}\} - \mathsf{E}_{H^{\boldsymbol{\lambda}_D}}\{\theta_{k,k}\} + \left(\mathsf{E}_{H^{\boldsymbol{\lambda}_D}}\{\theta_{k,k}\} - \mathsf{E}_{H^{\boldsymbol{\lambda}_O}}\{\theta_{k,l}\}\right)\left\{1 + (I-1)\mathsf{Pr}(\xi_i = \xi_j)\right\}. \tag{18}$$

Further, the equivalent expressions (16), (17) and (18) are proof of the following not so evident equalities

$$\mathsf{E}\{m_1\} = \frac{1}{I}\mathsf{E}\left\{\sum_{l=1}^{L} m_l^2\right\} = 1 + (I-1)\mathsf{Pr}(\xi_i = \xi_j).$$

We are unaware of any proof of this result in the literature on Poisson-Dirichlet processes.

In the case of $\bar{\kappa} = \mathsf{Var}(D_i|\boldsymbol{\lambda}, \boldsymbol{\varsigma})$, we used similar arguments and the fact that y_{ij} and y_{ih} are conditionally independent to derive the second moment first,

$$\mathsf{E}\{D_i^2|\boldsymbol{\lambda}, \boldsymbol{\varsigma}\} = \mathsf{E}\left\{\left(\sum_{j\neq i} y_{ij}\right)^2 \middle| \boldsymbol{\lambda}, \boldsymbol{\varsigma}\right\}$$
$$= [I-1]\mathsf{E}\left\{\mathsf{E}(y_{ij}^2|\xi_i, \xi_j, \boldsymbol{\Theta})\right\} + \frac{[I-1][I-2]}{2}\mathsf{E}\left\{\mathsf{E}(y_{ij}|\xi_i, \xi_j, \boldsymbol{\Theta})\mathsf{E}(y_{ih}|\xi_i, \xi_h, \boldsymbol{\Theta})\right\}.$$

Now $\mathsf{E}\left\{\mathsf{E}(y_{ij}^2|\xi_i, \xi_j, \boldsymbol{\Theta})\right\} = \mathsf{Pr}(\xi_i = \xi_j)\mathsf{E}_{H^{\boldsymbol{\lambda}_D}}\{\theta_{k,k}\} + \mathsf{Pr}(\xi_i \neq \xi_j)\mathsf{E}_{H^{\boldsymbol{\lambda}_O}}\{\theta_{k,l}\}$ and

$$\mathsf{E}\left\{\mathsf{E}(y_{ij}|\xi_i, \xi_j, \boldsymbol{\Theta})\mathsf{E}(y_{ih}|\xi_i, \xi_h, \boldsymbol{\Theta})\right\} = \mathsf{E}\{\theta_{k,k}^2\}\mathsf{Pr}(\xi_i = \xi_j = \xi_h) + \mathsf{E}\{\theta_{k,k}\theta_{k,l}\}\mathsf{Pr}(\xi_i = \xi_j \neq \xi_h) +$$
$$\mathsf{E}\{\theta_{k,k}\theta_{k,l}\}\mathsf{Pr}(\xi_i = \xi_h \neq \xi_j) + \mathsf{E}\{\theta_{k,l}^2\}\mathsf{Pr}(\xi_i \neq \xi_j = \xi_h) +$$
$$\left(\mathsf{E}\{\theta_{k,l}\}\right)^2\mathsf{Pr}(\xi_i \neq \xi_j \neq \xi_h).$$

Again, we can use the Pólya urn representation of the Poisson Dirichlet process to show that $\mathsf{Pr}(\xi_i = \xi_j \neq \xi_h) = \mathsf{Pr}(\xi_i = \xi_h \neq \xi_j) = \mathsf{Pr}(\xi_i \neq \xi_j = \xi_h) = \frac{(1-\alpha)(\beta+\alpha)}{(\beta+1)(\beta+2)}$, $\mathsf{Pr}(\xi_i \neq \xi_j \neq \xi_h) = \frac{(\beta+\alpha)(\beta+2\alpha)}{(\beta+1)(\beta+2)}$ and $\mathsf{Pr}(\xi_i = \xi_j = \xi_h) = \frac{(1-\alpha)(2-\alpha)}{(\beta+1)(\beta+2)}$.

Finally, $\bar{\Delta}$ follows again from a simple conditioning argument. Note that the conditional independence of y_{ij} and y_{ji} implies that

$$\Delta^N = \mathsf{Cov}(y_{ij}, y_{ji}) = \mathsf{Cov}\left\{\mathsf{E}(y_{ij}|\xi_i, \xi_j, \boldsymbol{\Theta}), \mathsf{E}(y_{ji}|\xi_i, \xi_j, \boldsymbol{\Theta})\right\}$$

Now

$$\mathsf{E}\left\{\mathsf{E}(y_{ij}|\xi_i, \xi_j, \boldsymbol{\Theta})\mathsf{E}(y_{ji}|\xi_i, \xi_j, \boldsymbol{\Theta})\right\} = \mathsf{E}_{H^{\boldsymbol{\lambda}_D}}\{\theta_{k,k}^2\}\mathsf{Pr}(\xi_i = \xi_j) + \mathsf{E}_{H^{\boldsymbol{\lambda}_O}}\{\theta_{k,l}\theta_{l,k}\}\mathsf{Pr}(\xi_i \neq \xi_j)$$

and,

$$\mathsf{E}\left\{\mathsf{E}(y_{ij}|\xi_i, \xi_j, \boldsymbol{\Theta})\right\} = \mathsf{E}\left\{\mathsf{E}(y_{ji}|\xi_i, \xi_j, \boldsymbol{\Theta})\right\} = \mathsf{E}_{H^{\boldsymbol{\lambda}_D}}\{\theta_{k,k}\}\mathsf{Pr}(\xi_i = \xi_j) + \mathsf{E}_{H^{\boldsymbol{\lambda}_O}}\{\theta_{k,l}\}\mathsf{Pr}(\xi_i \neq \xi_j).$$

Hence,

$$\begin{aligned}
\Delta^N &= \mathsf{E}\left\{\mathsf{E}(y_{ij}|\xi_i, \xi_j, \boldsymbol{\Theta})\mathsf{E}(y_{ji}|\xi_i, \xi_j, \boldsymbol{\Theta})\right\} - \mathsf{E}\left\{\mathsf{E}(y_{ij}|\xi_i, \xi_j, \boldsymbol{\Theta})\right\}\mathsf{E}\left\{\mathsf{E}(y_{ji}|\xi_i, \xi_j, \boldsymbol{\Theta})\right\} \\
&= \mathsf{Pr}(\xi_i = \xi_j)\mathsf{E}_{H^{\boldsymbol{\lambda}_D}}\{\theta_{k,k}^2\} + \mathsf{Pr}(\xi_i \neq \xi_j)\left(\mathsf{E}_{H^{\boldsymbol{\lambda}_O}}\{\theta_{k,l}\}\right)^2 \\
&\qquad - \left[\mathsf{Pr}(\xi_i = \xi_j)\mathsf{E}_{H^{\boldsymbol{\lambda}_D}}\{\theta_{k,k}\} + \mathsf{Pr}(\xi_i \neq \xi_j)\mathsf{E}_{H^{\boldsymbol{\lambda}_O}}\{\theta_{k,l}\}\right]^2.
\end{aligned}$$

Note that $\mathsf{E}\left\{\mathsf{E}(y_{ij}^2|\xi_i, \xi_j, \boldsymbol{\Theta})\right\} = \mathsf{E}\left\{\mathsf{E}(y_{ji}^2|\xi_i, \xi_j, \boldsymbol{\Theta})\right\}$, therefore $\mathsf{Var}(y_{ij}|\boldsymbol{\lambda}, \boldsymbol{\varsigma}) = \mathsf{Var}(y_{ji}|\boldsymbol{\lambda}, \boldsymbol{\varsigma})$ and

$$\begin{aligned}
\Delta^D &= \mathsf{Var}(y_{ij}|\boldsymbol{\lambda}, \boldsymbol{\varsigma}) \\
&= \mathsf{E}\left\{\mathsf{E}(y_{ij}^2|\xi_i, \xi_j, \boldsymbol{\Theta})\right\} - \left[\mathsf{E}\left\{\mathsf{E}(y_{ij}|\xi_i, \xi_j, \boldsymbol{\Theta})\right\}\right]^2 \\
&= \mathsf{Pr}(\xi_i = \xi_j)\mathsf{E}_{H^{\boldsymbol{\lambda}_D}}\{\theta_{k,k}\} + \mathsf{Pr}(\xi_i \neq \xi_j)\mathsf{E}_{H^{\boldsymbol{\lambda}_O}}\{\theta_{k,l}\} \\
&\qquad - \left[\mathsf{Pr}(\xi_i = \xi_j)\mathsf{E}_{H^{\boldsymbol{\lambda}_D}}\{\theta_{k,k}\} + \mathsf{Pr}(\xi_i \neq \xi_j)\mathsf{E}_{H^{\boldsymbol{\lambda}_O}}\{\theta_{k,l}\}\right]^2 \\
&= \bar{\theta}\{1 - \bar{\theta}\}.
\end{aligned}$$

D Moment generating function for D_i

Since observations are exchangeable we compute the moment generating function for D_1, without lost of generality assume that subject 1 is in the first community. Hence,

$$\begin{aligned}
\mathsf{E}\left(\exp\left\{tD_1\right\}\right) &= \mathsf{E}\left(\exp\left\{t\sum_{j\neq 1} y_{1,j}\right\}\right) \\
&= \mathsf{E}\left(\mathsf{E}\left[\exp\left\{t\sum_{j\neq 1} y_{1,j}\right\}\middle|\boldsymbol{\Theta},\xi_1,\ldots,\xi_I\right]\right) \\
&= \mathsf{E}\left(\mathsf{E}\left[\left\{\theta_{1,1}(\exp\left\{t\right\}-1)+1\right\}^{m_1-1}\prod_{k=2}^{K}\left\{\theta_{1,k}(\exp\left\{t\right\}-1)+1\right\}^{m_k}\middle|\boldsymbol{\Theta}\right]\right) \\
&= \mathsf{E}\left\{\left[1+\sum_{s=1}^{m_1-1}\binom{m_1-1}{s}\mathsf{E}_{H^{\boldsymbol{\lambda}_D}}(\theta_{1,1}^s)(\exp\{t\}-1)^s\right]\right. \\
&\qquad\left.\times\prod_{k=2}^{K}\left[1+\sum_{s=1}^{m_k}\binom{m_k}{s}\mathsf{E}_{H^{\boldsymbol{\lambda}_O}}(\theta_{1,k}^s)(\exp\{t\}-1)^s\right]\right\}
\end{aligned}$$

Note that, even though there is no close-form solution for the outward expectation, the expression can still be used to compute the moments of D_i by relating its moments to those of $m_1,\ldots,m_K$. Indeed, taking the first derivative with respect to t and evaluating at $t=0$ we get

$$\mathsf{E}(D_1|\boldsymbol{\lambda},\boldsymbol{\varsigma}) = I\mathsf{E}_{H^{\boldsymbol{\lambda}_O}}\{\theta_{1,l}\} - \mathsf{E}_{H^{\boldsymbol{\lambda}_D}}\{\theta_{1,1}\} + \left(\mathsf{E}_{H^{\boldsymbol{\lambda}_D}}\{\theta_{1,1}\} - \mathsf{E}_{H^{\boldsymbol{\lambda}_O}}\{\theta_{1,l}\}\right)\mathsf{E}\{m_1\},$$

as discussed in Appendix C.

E Proof of lemma 4

Note that

$$\begin{aligned}
\mathsf{Pr}(y_{ij}=1,y_{ih}=1,y_{jh}=1|\boldsymbol{\lambda},\boldsymbol{\varsigma}) &= \mathsf{E}\left\{\mathsf{Pr}(y_{ij}=1,y_{ih}=1,y_{jh}=1|\boldsymbol{\Theta},\xi_i,\xi_j,\xi_h,\boldsymbol{\lambda},\boldsymbol{\varsigma})\right\} \\
&= \mathsf{E}\{\theta_{k,k}^3\}\mathsf{Pr}(\xi_i=\xi_j=\xi_h) + \mathsf{E}\{\theta_{k,k}\theta_{k,l}^2\}\mathsf{Pr}(\xi_i=\xi_j\neq\xi_h) + \\
&\quad \mathsf{E}\{\theta_{k,k}\theta_{k,l}^2\}\mathsf{Pr}(\xi_j=\xi_h\neq\xi_i) + \mathsf{E}\{\theta_{k,k}\theta_{k,l}^2\}\mathsf{Pr}(\xi_i=\xi_h\neq\xi_j) + \\
&\quad \left(\mathsf{E}\{\theta_{k,l}\}\right)^3\mathsf{Pr}(\xi_i\neq\xi_j\neq\xi_h).
\end{aligned}$$

Again, because of the exchangeability, we can use the Pólya urn representation of the process to obtain $\mathsf{Pr}(\xi_i=\xi_j\neq\xi_h) = \mathsf{Pr}(\xi_j=\xi_h\neq\xi_i) = \mathsf{Pr}(\xi_i=\xi_h\neq\xi_j) = \frac{(1-\alpha)(\beta+\alpha)}{(\beta+1)(\beta+2)}$, $\mathsf{Pr}(\xi_i=\xi_j=\xi_h) = \frac{(1-\alpha)(2-\alpha)}{(\beta+1)(\beta+2)}$ and $\mathsf{Pr}(\xi_i\neq\xi_j\neq\xi_h) = \frac{(\beta+\alpha)(\beta+2\alpha)}{(\beta+1)(\beta+2)}$.

F Details on the MCMC sampler

F.1 Split-Merge MCMC Algorithm

1. Uniformly at random select two networks a and b.

2A. If $\boldsymbol{\zeta}_a = \boldsymbol{\zeta}_b$ propose a SPLIT move:

2A.1 Let $S.ab = \{j : \boldsymbol{\zeta}_j = \boldsymbol{\zeta}_a = \boldsymbol{\zeta}_b\}$, $S.a = \{a\}$ and $S.b = \{b\}$.
Assign the rest of $S.ab$'s components to either $S.a$ or $S.b$ at random with equal probability.

2A.2 Generate ξ_a and ξ_b through a modified Poyla urn scheme.

2A.3 For ξ_{ab} use the actual value of $\xi_{\boldsymbol{\zeta}_a}$

2A.4 Run a Gibbs sampler to update ξ_a , ξ_b and ξ_{ab}.

2B. Otherwise, if $\boldsymbol{\zeta}_a \neq \boldsymbol{\zeta}_b$ propose a MERGE move:

2B.1 Let $S.a = \{j : \boldsymbol{\zeta}_j = \boldsymbol{\zeta}_a\}$, $S.b = \{j : \boldsymbol{\zeta}_j = \boldsymbol{\zeta}_b\}$ and $S.ab = S.a \cup S.b$

2B.2 Generate ξ_{ab} through a modified Poyla urn scheme.

2B.3 For ξ_a and ξ_b use the actual value of $\xi_{\boldsymbol{\zeta}_a}$ and $\xi_{\boldsymbol{\zeta}_b}$, respectively.

2B.4 Run a Gibbs sampler to update ξ_a , ξ_b and ξ_{ab}.

3. Evaluate the proposal using a Metropolis-Hastings acceptance ratio. If the proposal is accepted, then $\boldsymbol{\zeta}$ and $\boldsymbol{\xi}$ change.

4. Each $\boldsymbol{\xi}_r$ is updated using a regular Gibbs sampler, regardless of the result of the M-H evaluation.

F.2 Modified Poyla urn scheme

Initialize each $\boldsymbol{\xi}_r$, $r = 1, \ldots, R$ as follows, let $\boldsymbol{\pi}$ be a random permutation of $\{1, \ldots, I\}$, set $\xi_{r,\pi_1} = 1$, next $\xi_{r,\pi_2} = 1$ with probability $(1 - \alpha_{2,r})/(\beta_{2,r} + 1)$, $\xi_{r,\pi_2} = 2$ otherwise. Then, for $h = 3, \ldots, I$, $\xi_{r,\pi_h} = k$ with probability

$$p\left(\xi_{r,\pi_h} = k \;\middle|\; \{\mathbf{Y}_j^{(h)} : \zeta_j = r\}\right)$$
$$= \begin{cases} \frac{m_{r,k} - \alpha_r}{\beta_r + h - 1} \prod_{\{j:\zeta_j = r\}} p\left(\mathbf{y}_{\pi_h,j}^{(h-1)} \;\middle|\; \{\mathbf{y}_{l,j}^{(h-1)} : l \in \boldsymbol{\pi}^{(h-1)}, \xi_{r,l} = k\}\right) & k \leq L_r^{(h-1)} \\ \frac{\beta_r + \alpha_r L_r^{(h-1)}}{\beta_r + h - 1} \prod_{\{j:\zeta_j = r\}} p\left(\mathbf{y}_{\pi_h,j}^{(h-1)}\right) & k = L_r^{(h-1)} + 1 \end{cases},$$

where $\mathbf{Y}_j^{(h)} = \{y_{i,i',j} : i, i' \in \boldsymbol{\pi}^{(h)}\}$, $\mathbf{y}_{l,j}^{(h-1)} = \{y_{i,i',j} : (i,i') \in \{i = l, i' \in \boldsymbol{\pi}^{(h-1)}\} \cup \{i \in \boldsymbol{\pi}^{(h-1)}, i' = l\}\}$, with $\boldsymbol{\pi}^{(h)} = \{\pi_1, \ldots, \pi_h\}$, $L_r^{(h-1)}$ is the total number of groups and $m_{r,k}^{(h-1)}$ is the number of actors in the k-th group after assigning the first $h-1$ actors. There,

$$
\begin{aligned}
& p\left(\mathbf{y}_{\pi_h,j}^{(h-1)} \;\middle|\; \{\mathbf{y}_{l,j}^{(h-1)} : l \in \boldsymbol{\pi}^{(h-1)}, \xi_{r,l} = k\}\right) \\
&= \begin{cases} p_\theta\left(\mathbf{y}_{(\pi_h),k,j}^{(h-1)} \;\middle|\; \mathbf{y}_{k,k,j}^{(h-1)}\right) \prod_{l=1, l\neq k}^{L_r^{(h-1)}} p_\theta\left(\mathbf{y}_{(\pi_h),l,j}^{(h-1)} \;\middle|\; \mathbf{y}_{k,l,j}^{(h-1)}\right) p_\theta\left(\mathbf{y}_{l,(\pi_h),j}^{(h-1)} \;\middle|\; \mathbf{y}_{l,k,j}^{(h-1)}\right) & \text{if } j \text{ is directed} \\ p_\theta\left(\mathbf{y}_{(\pi_h),k,j}^{(h-1)}, \mathbf{y}_{k,(\pi_h),j}^{(h-1)} \;\middle|\; \mathbf{y}_{k,k,j}^{(h-1)}\right) \prod_{l=1, l\neq k}^{L_r^{(h-1)}} p_\theta\left(\mathbf{y}_{(\pi_h),l,j}^{(h-1)}, \mathbf{y}_{l,(\pi_h),j}^{(h-1)} \;\middle|\; \mathbf{y}_{k,l,j}^{(h-1)}, \mathbf{y}_{l,k,j}^{(h-1)}\right) & \text{if } j \text{ is undirected} \end{cases},
\end{aligned}
$$

where $\mathbf{y}_{(\pi_h),l,j}^{(h-1)} = \{y_{\pi_h,i',j} : i' \in \boldsymbol{\pi}^{(h-1)}; \xi_{r,i'}^{(h-1)} = l\}$, $\mathbf{y}_{l,(\pi_h),j}^{(h-1)} = \{y_{i,\pi_h,j} : i \in \boldsymbol{\pi}^{(h-1)}; \xi_{r,i}^{(h-1)} = l\}$, $\mathbf{y}_{k,l,j}^{(h-1)} = \{y_{i,i',j} : i, i' \in \boldsymbol{\pi}^{(h-1)}; \xi_{r,i}^{(h-1)} = k; \xi_{r,i'}^{(h-1)} = l\}$ and $p_\theta(\cdot \mid \cdot)$ can be found using $p_\theta(\cdot)$ the marginal posterior of $\mathbf{y}$ under a given prior for θ. Moreover,

$$
p\left(\mathbf{y}_{\pi_h,j}^{(h-1)}\right) = \begin{cases} \prod_{l=1}^{L_r^{(h-1)}} p_\theta\left(\mathbf{y}_{(\pi_h),l,j}^{(h-1)}\right) p_\theta\left(\mathbf{y}_{l,(\pi_h),j}^{(h-1)}\right) & \text{if } j \text{ is directed} \\ \prod_{l=1}^{L_r^{(h-1)}} p_\theta\left(\mathbf{y}_{(\pi_h),l,j}^{(h-1)}, \mathbf{y}_{l,(\pi_h),j}^{(h-1)}\right) & \text{if } j \text{ is undirected} \end{cases}.
$$

Let q_r be the product of the probabilities of the assignment made at each step. Hence, q_r is the probability of obtaining $\boldsymbol{\xi}_r$, which will be needed for the Metropolis-Hastings ratio.

F.3 Metropolis-Hastings ratio

For $\boldsymbol{\eta} = (\boldsymbol{\zeta}, \boldsymbol{\xi})$ if $\boldsymbol{\eta}^*$ is a SPLIT move, accept a SPLIT move with probability

$$
\begin{aligned}
& a\left(\boldsymbol{\eta}^* \mid \boldsymbol{\eta}\right) \\
&\quad = min\left\{ \frac{2^{n_{ab}-2}(\beta_1 + R\alpha_1)}{\Gamma(1-\alpha_1)} \frac{\Gamma(n_a - \alpha_1)\Gamma(n_b - \alpha_1)}{\Gamma(n_{ab} - \alpha_1)} \frac{p\left(\mathbf{Y}, \boldsymbol{\xi}_a; S.a\right) p\left(\mathbf{Y}, \boldsymbol{\xi}_b; S.b\right)}{p\left(\mathbf{Y}\boldsymbol{\xi}_{ab}; S.ab\right)} \frac{Q.ab}{Q.aQ.b}, 1 \right\},
\end{aligned}
$$

where $R =$ the total number of classes before the split, $n_r = |S.r|$ and

$$
p\left(\mathbf{Y}, \boldsymbol{\xi}_r; S.r\right) = \prod_{j \in S.r} p\left(\mathbf{Y}_j \mid \boldsymbol{\xi}_r\right) p(\boldsymbol{\xi}_r | \alpha_{2,r}, \beta_{2,r}),
$$

with $p(\boldsymbol{\xi}_r | \alpha_{2,r}, \beta_{2,r})$ given in equation (13) and

$$
p\left(\mathbf{Y}_j \mid \boldsymbol{\xi}_r\right) \;=\; \begin{cases} \prod_{l=1}^{L_r} p_\theta\left(\mathbf{y}_{l,l,j}\right) \prod_{l'>l}^{L_r} p_\theta\left(\mathbf{y}_{l,l',j}\right) p_\theta\left(\mathbf{y}_{l',l,j}\right) & \text{if } j \text{ is directed} \\ \prod_{l=1}^{L_r} p_\theta\left(\mathbf{y}_{l,l,j}\right) \prod_{l'>l}^{L_r} p_\theta\left(\mathbf{y}_{l,l',j}, \mathbf{y}_{l',l,j}\right) & \text{if } j \text{ is undirected} \end{cases},
$$

where $\mathbf{y}_{l,l',j} = \{y_{i,i',j} : \xi_{r,i} = l; \xi_{r,i'} = l'\}$ and $p_\theta(\cdot)$ the marginal posterior of $\mathbf{y}$ under a given prior for θ.

$Q.r = p(\alpha_{2,r})q(\xi_r^{(0)})q(\xi_r^{(1)})$ where $p(\alpha_{2,r})$ is the prior on $\alpha_{2,r}$, and $q(\xi_r^{(0)})$ and $q(\xi_r^{(1)})$ are the probabilities of obtaining $\xi_r^{(0)}$ and $\xi_r^{(1)}$ respectively, using the method described in section F.2.

References

AIROLDI, E. M., BLEI, D. M., FIENBERG, S. E. & XING, E. P. (2008). Mixed membership stochastic blockmodels. *Journal of Machine Learning Research* **9**, 1981–2014.

ALDOUS, D. J. (1981). Representations for partially exchangeable arrays of random variables. *Journal of Multivariate Analysis* **11**, 581–598.

BICKEL, P. J. & CHEN, A. (2009). A nonparametric view of network models and Newman girvan and other modularities. *Proceedings of the National Academy of Sciences* **106**, 21068–21073.

CLAUSET, A. (2005). Finding local community structure in networks. *Physical Review E* **72**, 026132.

DAHL, D. (2003). An improved merge-split sampler for conjugate Dirichlet process mixture models. Technical report, Department of Statistics, University of Winsconsin.

ERDÖS, P. & RÉNYI, A. (1959). On random graphs. *Publicationes Mathematicae* **6**, 290–297.

FIENBERG, S. E., MEYER, M. M. & WASSERMAN, S. S. (1985). Statistical analysis of multiple sociometric relations. *Journal of the American Statistical Association* **80**, 51–67.

FRANK, O. & STRAUSS, D. (1986). Markov graphs. *Journal of the American Statistical Association* **81**, 832–842.

GALASKIEWICZ, J. & MARSDEN, P. V. (1978). Interorganizational resource neworks: Formal patterns of averlap. *Social Science Research* **7**, 89–107.

GIRVAN, M. & NEWMAN, M. E. J. (2002). Community structure in social and biological networks. *Proceedings of the National Academy of Sciences* **99**, 7821–7826.

HANDCOCK, M. S., RAFTERY, A. E. & TANTRUM, J. M. (2007). Model-based clustering for social networks. *Journal of the Royal Statistical Society, Series A* **170**, 301–354.

HOFF, P. D., RAFTERY, A. E. & HANDCOCK, M. S. (2002). Latent space approaches to social network analysis. *Journal of American Statistical Association* **97**, 1090–1098.

HOLLAND, P. W. & LEINHARDT, S. (1975). Local structure in social networks. In *Sociological Methodology 1976*, Ed. D. R. Heise, pp. 1–54. San Francisco: Jossey-Bass.

JAIN, S. & NEAL, R. M. (2004). A split-merge Markov chain Monte Carlo procedure for the Dirichlet process mixture model. *Journal of Graphical and Computational Statistics* **13**, 158–182.

KEMP, C., TENENBAUM, J. B., GRIFFITHS, T. L., YAMADA, T. & UEDA, N. (2006). Learning systems of concepts with an infinite relational model. In *Proceedings of the 22nd Annual Conference on Artificial Intelligence*.

KRACKHARDT, D. (1987). Cognitive social structures. *Social Networks* **9**, 104–134.

LAU, J. W. & GREEN, P. (2007). Bayesian model based clustering procedures. *Journal of Computational and Graphical Statistics* **16**, 526–558.

LORRAIN, F. P. & WHITE, H. C. (1971). Structural equivalence of individuals in social networks. *Journal of Mathematical Sociology* **1**, 49–80.

MISHRA, N., SCHREIBER, R., STANTON, I. & TARJAN, R. E. (2008). Finding strongly knit clusters in social networks. *Internet Mathematics* **5**, 155–174.

NEWMAN, M. E. J. (2003a). Mixing patterns in networks. *Physical Review E* **67**, 026126.

NEWMAN, M. E. J. (2003b). The structure and function of complex networks. *SIAM Review* **45**, 167–256.

NEWMAN, M. E. J. (2004). Detecting community structure in networks. *European Physics Journal B* **38**, 321–330.

PITMAN, J. (1995). Exchangeable and partially exchangeable random partitions. *Probability Theory and Related Fields* **102**, 145–158.

RODRIGUEZ, A., REYES, P. E. & VUPPALA, K. (2011). Nonparametric bayesian modeling under partial exchangeability: A review. Technical report, University of California - Santa Cruz.

ROETHLISBERGER, F. & DICKSON, W. (1939). *Management and the worker*. Cambridge: Cambridge University Press.

SETHURAMAN, J. (1994). A constructive definition of Dirichlet priors. *Statistica Sinica* **4**, 639–650.

WANG, Y. J. & WONG, G. Y. (1987). Stochastic blockmodels for directed graphs. *Journal of the American Statistical Association* **82**, 8–19.

WASSERMAN, S. S. & FAUST, K. (1994). *Social Network Analysis: Methods and Applications*. Cambridge: Cambridge University Press.

WHITE, H. C., BOORMAN, S. A. & BREIGER, R. L. (1976). Social structure from multiple networks: I, Blockmodels of roles and positions. *American Journal of Sociology* **81**, 730–780.

XU, Z., TRESP, V., YU, K. & KRIEGEL, H.-P. (2006). Infinite hidden relational models. In *Proceedings of the 22nd Annual Conference on Uncertainty in Artificial Intelligence*.